arXiv:math/0202285v1 [math.GR] 27 Feb 2002

STALLINGS FOLDINGS AND SUBGROUPS OF FREE GROUPS

ILYA KAPOVICH AND ALEXEI MYASNIKOV

Abstract. We re-cast in a more combinatorial and computational form the topological approach of J.Stallings to the study of subgroups of free groups.

1. Introduction

The subgroup structure of free groups is a classical subject which goes back to the origins of Group Theory. The original approach developed by Nielsen was to treat this topic combinatorially, mainly using the technique of Nielsen transformations. Indeed, even to this day this method remains among the most powerful ones for working with subgroups of free groups. The development of algebraic topology and covering space theory in the 1940s suggested a different, more geometric approach. A free group F can be identified with the fundamental group of a topological graph (which we may think of as a 1-complex). Then any subgroup of F corresponds to a covering map from another graph to the original graph. The topological viewpoint was studied in detail by J.Stallings in a seminal paper [43]. In this work J.Stallings introduced an extremely useful notion of a *folding* of graphs. The ideas of Stallings have found many interesting applications (see, for example, [15], [16], [18], [44],,[45], [3], [41], [11], [13], [14], [48], [49], [50]). However, most of these applications, as far as the theory of free groups is concerned, were for the study of automorphisms of free groups. Thus, for example, the notion of a folding plays an important role in the construction of train-track maps due to M.Bestvina and M.Handel [6], [7].

In this paper we want to recast the ideas of Stallings' work [43] in a more combinatorial guise and apply them more systematically to the subgroup structure of free groups.

Namely, to every subgroup H of a free group $F(X)$ we will associate a directed graph $\Gamma(H)$ whose edges are labeled by the elements of X. It turns out that this graph $\Gamma(H)$ (which is finite and very easy to construct if H is finitely generated) carries all the essential information about the subgroup H itself. Geometrically, the graph $\Gamma(H)$ represents the topological core of the covering space, corresponding to H, of the wedge of $\#X$ circles. Algebraically, $\Gamma(H)$ can also be viewed as the "essential part" of the relative coset Cayley graph of G/H with respect to X. This last approach was used by C.Sims in [42]. Yet another view of $\Gamma(H)$ comes from Bass-Serre theory. Namely, $F(X)$ acts on its standard Cayley graph, which is a regular tree. There is a unique H-invariant subtree $T(H)$ which is minimal among H-invariant subtrees containing 1. Our $\Gamma(H)$ can be identified with the quotient $T(H)/H$.

However, in the present paper we will consider and study $\Gamma(H)$ almost exclusively as a combinatorial object. Namely, we will view it as a labeled directed graph and, also, as an automaton over the alphabet $X \cup X^{-1}$. There are several reasons why we believe this approach is interesting and worthwhile.

First and foremost, the graph-automaton $\Gamma(H)$ turns out to be an extremely useful and natural object when one considers various algorithmic and computational problems for free groups and their subgroups (e.g. the membership problem). In fact, most classical algorithms dealing with such problems usually use Nielsen methods and involve enumerating all elements of $F(X)$ of bounded length. Since $F(X)$ has exponential growth, all such algorithms require at least an exponential amount of time. However, it turns out that looking at $\Gamma(H)$ allows one to solve most of these problem much faster, ordinarily in polynomial time. In fact, much of this work is motivated by the desire to describe various algorithms implemented in the Computational Group Theory Software Package MAGNUS. It is also important that

---

1991 *Mathematics Subject Classification.* 20F.

the graph-automaton $\Gamma(H)$ is a canonical object which gives us some advantages as compared to the standard situation in the theory of automata and formal languages.

Second, we thought it worth-while to give an account of Stallings' ideas in a manner accessible to people with little topological or even group-theoretic background. There have been a number of papers where these ideas were applied to some specific subgroup properties of free groups. However, we wanted to give a self-contained, elementary and comprehensive treatment of the subject. For this reason we included complete and independent proofs of most basic facts, a substantial number of explicit examples and a wide assortment of possible applications.

Thus we re-prove many classical, well-known or folklore results about the subgroup structure of free groups. For instance, we prove the Marshal Hall Theorem and the Greenberg-Stallings Theorem as well as establish some well-known facts regarding ascending and descending chains of subgroups in free groups. In some instances we choose not to prove certain statements in the strongest possible form, as our real goal is to demonstrate the usefulness of the method.

However, there are quite a few results that appear to be new, or at least, have never been published. Thus in Section 11 we develop a mini-theory of "algebraic" and "free" extensions in the context of free groups. This technique is then applied in Section 13 to study malnormal closures and isolators of finitely generated subgroups of free groups.

This paper is based, in large measure, on the notes of a Group Theory Research Seminar, that has been running at the CUNY Graduate Center in the Spring and Fall semesters of 1999. Our special thanks go to Gilbert Baumslag, Toshiaki Jitsukawa, Gretchen Ostheimer, Gillian Elston, Robert Gilman, Vladimir Shpilrain, Sean Cleary, Lev Shneyerson, Fuh Ching-Fen, Dmitri Pechkin, Dmitri Bormotov, Alexei D. Myasnikov, Alexei Kvaschuk, Denis Serbin Katalin Bencsath, Arthur Sternberg and other regular participants of the seminar for the many lively, engaging and stimulating discussions.

2. Labeled graphs

Definition 2.1 (X-digraph)**.** Let $X = \{x_1, \dots, x_N\}$ be a finite alphabet. By a *X-labeled directed graph* Γ (also called a *X-digraph* or even just *X-graph*) we mean the following:

Γ is a combinatorial graph where every edge e has an arrow (direction) and is labeled by a letter from X, denoted $\mu(e)$. (Note that Γ can be either finite or infinite.)

For each edge e of Γ we denote the origin of e by $o(e)$ and the terminus of e by $t(e)$. If $o(e) = t(e)$ then e is a *loop*.

There is an obvious notion of a *morphism* between two X-digraphs. Namely, if Γ and Δ are X-digraphs then a map $\pi : \Gamma \longrightarrow \Delta$ is called a *morphism* of X-digraphs, if π takes vertices to vertices, directed edges to directed edges, preserves labels of directed edges and has the property that $o(\pi(e)) = \pi(o(e)), t(\pi(e)) = \pi(t(e))$ for any edge e of Γ.

Examples of X-graphs, with $X = \{a, b, c\}$ are shown in Figure 1.

For the remainder of the paper, unless specified otherwise, X will stand for a finite alphabet $X = \{x_1, \dots, x_N\}$. We shall also denote $\Sigma = X \cup X^{-1}$.

Convention 2.2. Given a X-digraph Γ, we can make Γ into an oriented graph labeled by the alphabet $\Sigma = X \cup X^{-1}$. Namely, for each edge e of Γ we introduce a formal inverse e^{-1} of e with label $\mu(e)^{-1}$ and the endpoints defined as $o(e^{-1}) = t(e)$, $t(e^{-1}) = o(e)$. The arrow on e^{-1} points from the terminus of e to the origin of e. For the new edges e^{-1} we set $(e^{-1})^{-1} = e$.

The new graph, endowed with this additional structure, will be denoted by $\hat{\Gamma}$. In fact in many instances we will abuse notation by disregarding the difference between Γ and $\hat{\Gamma}$.

The edge-set of $\hat{\Gamma}$ is naturally partitioned as $E\hat{\Gamma} = E\Gamma \cup \bar{E}\Gamma$. We will say that the edges e of Γ are *positively oriented* or *positive* in $\hat{\Gamma}$ and that their formal inverses e^{-1} are *negatively oriented* or *negative* in $\hat{\Gamma}$. We will also denote $E^{+}\hat{\Gamma} = E\Gamma$ and $E^{-}\hat{\Gamma} = (E\Gamma)^{-1}$.

The use of $\hat{\Gamma}$ allows us to define the notion of a *path* in Γ. Namely, a *path* p in Γ is a sequence of edges $p = e_1, \dots, e_k$ where each e_i is an edge of $\hat{\Gamma}$ and the origin of each e_i (for $i > 1$) is the terminus of e_{i-1}. In this situation we will say that the *origin* $o(p)$ of p is $o(e_1)$ and the *terminus* $t(p)$ is $t(e_k)$. The *length*

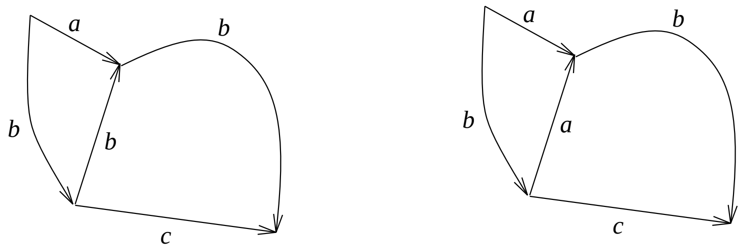

(a) (b)

Figure 1. Folded and non-folded graphs. The labeling alphabet is $X = \{a, b, c\}$.

$|p|$ of this path is set to be k. Also, such a path p has a naturally defined label $\mu(p) = \mu(e_1)\dots\mu(e_k)$. Thus $\mu(p)$ is a word in the alphabet $\Sigma = X \cup X^{-1}$. Note that it is possible that $\mu(p)$ contains subwords of the form aa^{-1} or $a^{-1}a$ for some $a \in X$.

Also, if v is a vertex of Γ, we will consider the sequence $p = v$ to be a path with $o(p) = t(p) = v$, $|p| = 0$ and $\mu(p) = 1$ (the empty word).

Definition 2.3 (Folded graphs)**.** Let Γ be an X-digraph. We say that Γ is *folded* if for each vertex v of Γ and each letter $a \in X$ there is at most one edge in Γ with origin v and label a and there is at most one edge with terminus v and label a.

The graph shown in Figure 1(a) is folded and the graph shown in Figure 1(b) is not folded. Note that Γ is folded if and only if for each vertex v of Γ and each $x \in \Sigma = X \cup X^{-1}$ there is at most one edge in $\hat{\Gamma}$ with origin v and label x. Thus in a folded X-digraph the degree of each vertex is at most $2\#(X)$.

The following notion of graph folding plays a fundamental role in this paper. Suppose Γ is an X-digraph and e_1, e_2 are edges of Γ with common origin and the same label $x \in \Sigma$. Then, informally speaking, *folding* Γ at e_1, e_2 means identifying e_1 and e_2 into a single new edge labeled x. The resulting graph carries a natural structure of an X-digraph. A more precise definition is given below.

Definition 2.4 (Folding of graphs)**.** Let Γ be an X-digraph. Suppose that v_0 is a vertex of Γ and f_1, f_2 are two distinct edges of $\hat{\Gamma}$ with origin v_0 and such that $\mu(f_1) = \mu(f_2) = x \in \Sigma = X \cup X^{-1}$ (so that Γ is not folded). Let h_i be the positive edge of Γ corresponding to f_i (that is $h_i = f_i$ if f_i is positive and $h_i = {f_i}^{-1}$ if f_i is negative). Note that, depending on whether $x \in X$ or $x \in X^{-1}$, the edges f_1 and f_2 are either both positive or both negative in $\hat{\Gamma}$.

Let Δ be an X-digraph defined as follows.

The vertex set of Δ is the vertex set of Γ with $t(f_1)$ and $t(f_2)$ removed and a new vertex t_f added (we think of the vertices $t(f_1)$ and $t(f_2)$ as being identified to produce vertex t_f):

$$V\Delta = \big(V\Gamma - \{t(f_1), t(f_2)\}\big) \cup \{t_f\}$$

The edge set of Δ is the edge set of Γ with the edges h_1, h_2 removed and a new edge h added (we think of the edges h_1 and h_2 as being identified or "folded" to produce a new edge h):

$$E\Delta = (E\Gamma - \{h_1, h_2\}) \cup \{h\}$$

The endpoints and arrows for the edges of Δ are defined in a natural way. Namely, if $e \in E\Delta$ and $e \neq h$ (that is $e \in E\Gamma$, $e \neq h_i$) then

1. we put $o_\Delta(e) = o_\Gamma(e)$ if $o_\Gamma(e) \neq t(f_i)$ and $o_\Delta(e) = t_f$ if $o_\Gamma(e) = t(f_i)$ for some i;
2. we put $t_\Delta(e) = t_\Gamma(e)$ if $t_\Gamma(e) \neq t(f_i)$ and $t_\Delta(e) = t_f$ if $t_\Gamma(e) = t(f_i)$ for some i

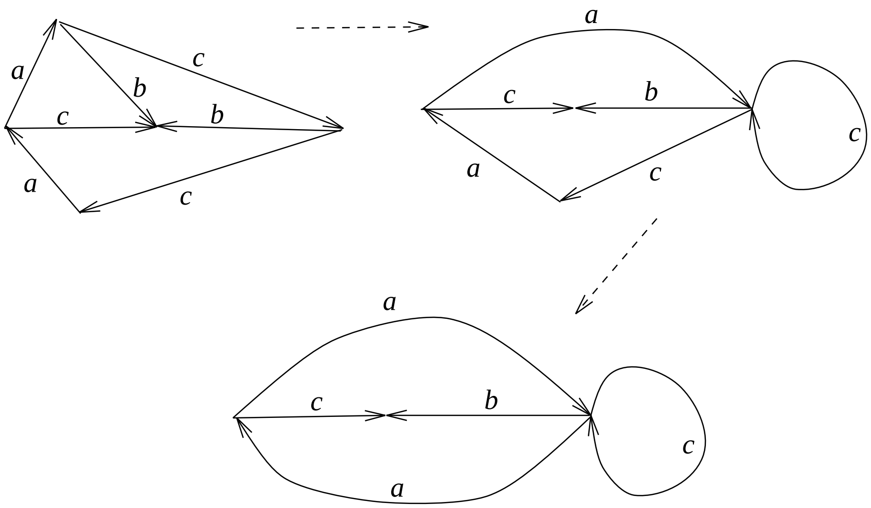

FIGURE 2. Folding of labeled graphs

For the edge h we put $o_\Delta(h) = o_\Gamma(h)$ if $o_\Gamma(h) \neq t(f_1), t(f_2)$ and $o_\Delta(h) = t_f$ if $o_\Gamma(h) = t(f_i)$ for some i. Similarly, we put $t_\Delta(h) = t_\Gamma(h)$ if $t_\Gamma(h) \neq t(f_1), t(f_2)$ and $t_\Delta(h) = t_f$ if $t_\Gamma(h) = t(f_i)$ for some i.

We define labels on the edges of Δ as follows: $\mu_\Delta(e) = \mu_\Gamma(e)$ if $e \neq h$ and $\mu_\Delta(h) = \mu_\Gamma(h_1) = \mu_\Gamma(h_2)$.

Thus Δ is an X-digraph. In this situation we say that Δ is obtained from Γ by a *folding* (or by *folding the edges* f_1 and f_2).

We suggest that the reader investigates carefully what happens in the above definition if one or both of the edges f_1, f_2 are loops. The notion of folding is illustrated in Figure 2.

We also need the following obvious statement summarizing some basic properties of folding.

Lemma 2.5. *Let Γ_1 be an X-digraph obtained by a folding from a graph Γ. Let v be a vertex of Γ and let v_1 be the corresponding vertex of Γ_1.*

Then the following holds:

1. *If Γ is connected then Γ_1 is connected.*
2. *Let p be a path from v to v in Γ with label w. Then the edge-wise image of p in Γ_1 is a path from v_1 to v_1 with label w.*
3. *Note that if an X-digraph Γ is finite, then a folding always decreases the number of edges in Γ by one.*

Convention 2.6 (Reduced words and reduced paths)**.** Recall that a word w in the alphabet $\Sigma = X \cup X^{-1}$ is said to be *freely reduced* if it does not contain a subword of the form aa^{-1} or $a^{-1}a$ for $a \in X$.

A path p in an X-digraph Γ is said to be *reduced* if p does not contain subpaths of the form e, e^{-1} for $e \in E\hat{\Gamma}$.

For an alphabet B we denote by B^* the set of all words (including the empty word 1) in the alphabet B.

Definition 2.7 (Language recognized by a digraph)**.** Let Γ be an X-digraph and let v be a vertex of Γ. We define *the language of* Γ *with respect to* v to be:

$$L(\Gamma, v) = \{\mu(p) \,|\, p \text{ is a reduced path in } \Gamma \text{ from } v \text{ to } v\}$$

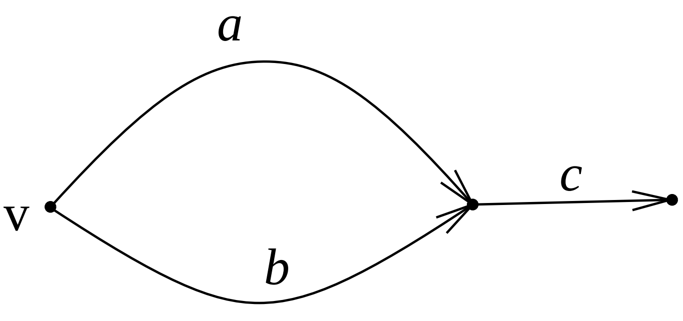

FIGURE 3. Language of a digraph. In this case the language $L(\Gamma, v)$ consists of all words of the form $(ab^{-1})^n$ where n is any integer (possibly negative). Note that the edge labeled c is irrelevant since there are no reduced paths from v to v passing through this edge.

If a word w belongs to $L(\Gamma, v)$, we will also sometimes say that w is *accepted* by (Γ, v) (or just by Γ if the choice of v is fixed).

Thus $L(\Gamma, v)$ is a subset of Σ^* where $\Sigma = X \cup X^{-1}$. An example of the language corresponding to a digraph is shown in Figure 3

Note that the words in $L(\Gamma, v)$ are not necessarily freely reduced. However, if the graph Γ is folded, then all the words in $L(\Gamma, v)$ are freely reduced.

Remark 2.8 (Labeled graphs versus standard automata)**.** It is possible and often useful to view the pair $(\hat{\Gamma}, v)$ as an automaton M (possibly infinite) over the alphabet $\Sigma = X \cup X^{-1}$. The vertices of Γ are states and the edges represent transitions. There is only one initial state, namely v and only one accept state, also v. Then Γ is folded if and only if the corresponding automaton M is deterministic. It should be noted, that our definition of the language $L(\Gamma, v)$ is slightly different from the standard notion of the language $L(M)$ recognized by M. Namely, using the standard conventions of the automata theory, the language of $L(M)$ would consist of the labels of all (not just reduced) paths from v to v in Γ. However, the set of all freely reduced words in Σ is easily seen to be a regular language (i.e. it is recognized by a finite state automaton). Since intersections of regular languages are regular, this implies that for a finite and folded X-digraph Γ, the language $L(\Gamma, v)$ is regular.

We formulate this obvious but important statement as a lemma.

Lemma 2.9. *Let Γ be a folded X-digraph. Let v be a vertex of Γ. Then all the words in the language $L(\Gamma, v)$ are freely reduced.*

3. LABELED GRAPHS AND SUBGROUPS OF FREE GROUPS

In this section we will explain how to every finitely generated subgroup of a free group one can associate a (reasonably) canonical labeled graph.

Let X be a finite alphabet and let $\Sigma = X \cup X^{-1}$. If w is a word in Σ, we will denote by $\overline{w}$ the freely reduced word in Σ obtained from w by performing all possible (if any) free reductions.

Recall that a *free group on X*, denoted $F(X)$ is the collection of all freely reduced words in Σ (including the empty word 1). The multiplication in $F(X)$ is defined as

$$f \cdot g := \overline{fg}, \qquad \text{for any } f, g \in F(X)$$

(Although not quite trivial, it can be shown that thus defined, the multiplication makes $F(X)$ a group [32]). The number of elements in X is called *the rank of* $F(X)$ and X is referred to as *a free basis of* $F(X)$.

For any element $g \in F(X)$ we will denote by $|g|_X$ the length of the unique freely reduced Σ-word, representing g. If w is a word in Σ (which may or may not be freely reduced), we will denote by $|w|$ the length of w.

It turns out that the languages of directed X-graphs correspond to subgroups of $F(X)$.

Namely, the following simple but important statement holds.

Proposition 3.1. *Let Γ be an X-digraph and let v be a base-vertex of Γ. Then the set*

$$\overline{L} = \{\overline{w} \mid w \in L(\Gamma, v)\}$$

is a subgroup of $F(X)$.

Proof. Indeed, let $v_1, v_2 \in \overline{L}$. Then there are reduced paths p_1 and p_2 from v to v in Γ such that the label of p_i is w_i and $\overline{w_i} = v_i$.

The concatenation p_1p_2 is a path in Γ from v to v which may or may not be path-reduced. Let p be the reduced path obtained from p_1p_2 by making all possible path reductions. This means that the label $w = \mu(p)$ is obtained from the word w_1w_2 by performing several free reductions (even though $w = \mu(p)$ may not be freely reduced itself). Therefore $\overline{w} = \overline{w_1w_2} = v_1 \cdot v_2 \in F(X)$. On the other hand w is the label of a reduced path from v to v and therefore $w \in L(\Gamma, v) = L$ (by definition). Thus $v_1 \cdot v_2 \in \overline{L}$ and $\overline{L}$ is closed under multiplication.

It is easy to see that the inverse path $(p_1)^{-1}$ of p_1 has label w_1^{-1}. This implies that $\overline{L}$ is closed under taking inverses. Also, obviously $1 \in \overline{L}$.

Thus $\overline{L}$ is a subgroup of $F(X)$, as required. □

Lemma 3.2. *Suppose Γ is a folded X-digraph. Then $L(\Gamma, v) = L = \overline{L}$ is a subgroup of $F(X)$.*

Proof. By Lemma 2.9 the labels of reduced paths in Γ are already freely reduced. Therefore by Proposition 3.1 $L(\Gamma, v) = L = \overline{L}$ is a subgroup of $F(X)$. □

The following lemma shows that subgroups corresponding to finite labeled graphs are finitely generated.

Lemma 3.3. *Let Γ be a connected X-digraph and let v be a vertex of Γ. For each vertex $u \neq v$ of Γ choose a reduced path p_u in Γ from v to u. Put $p_v = v$, the path of length zero consisting just of vertex v. For each edge e of $\hat{\Gamma}$ (whether positive or negative) put $p_e = p_{o(e)}e(p_{o(e)})^{-1}$ so that p_e is a path in Γ from v to v. Denote $[e] = \overline{\mu(p_e)}$.*

Then the subgroup $H = \overline{L(\Gamma, v)}$ of $F(X)$ is generated by the set

$$S = \{[e] \mid \text{ where } e \text{ is a positive edge of } \Gamma\}$$

In particular, if Γ is finite, the subgroup H is finitely generated.

Proof. The path p_e (if it is not already reduced) can be transformed into a reduced path p'_e (from v to v) by a series of path-reductions. Clearly $\mu(p'_e) \in L(\Gamma, v)$, $\overline{\mu(p'_e)} = \overline{\mu(p_e)}$ and hence $[e] = \overline{\mu(p_e)} \in \overline{L(\Gamma, v)} = H$. Thus $S \subseteq H$ and $\langle S \rangle \leq H$.

It remains to prove that any element of H can be expressed as a product of elements of S and their inverses. Note that by definition $p_{e^{-1}} = (p_e)^{-1}$ and so $[e^{-1}] = [e]^{-1}$. Thus it suffices to show that any element of H can be expressed as a product of elements $[e]$, where e is an edge of $\hat{\Gamma}$.

Let $h \in H$, $h \neq 1$. Then there is a nontrivial reduced path p from v to v with label w such that $\overline{w} = h$. Let $p = e_1, \dots, e_k$, where e_i are edges of $\hat{X}$. Further let $v_1 = v$, $v_{k+1} = v$ and let v_i be the initial vertex of e_i (and therefore the terminal vertex of e_{i+1}). Consider now the path

$$p' = p_{e_1} \dots p_{e_k} = p_{v_1}e_1p_{v_2}^{-1}p_{v_2}e_2p_{v_3}^{-1} \dots p_{v_k}e_kp_{v_{k+1}}^{-1}$$

It is obvious that the path p' can be transformed by path reductions into the path p. Therefore $\overline{\mu(p')} = \overline{\mu(p)} = h$. On the other hand

$$\overline{\mu(p')} = \overline{p_{e_1}} \dots \overline{p_{e_k}} = [e_1] \dots [e_k] \in \langle S \rangle$$

Thus $H = \langle S \rangle$ as required. □

The following simple but important observation plays a crucial role in this paper.

Lemma 3.4. *Let Γ be an X-digraph and let Γ' be an X-digraph obtained from Γ by a single folding. Further let v be a vertex of Γ and let v' be the corresponding vertex of Γ'.*

Then $\overline{L(\Gamma, v)} = \overline{L(\Gamma', v')}$.

Proof. Suppose Γ' is obtained from Γ by folding two edges e_1, e_2 in $\hat{\Gamma}$ which have the same initial vertex u and the same label $x \in X$. The edges e_1, e_2 are folded into an edge e of Γ' labeled x and with origin u'.

Suppose p is a reduced path in Γ from v to v, so that $\mu(p) \in L(\Gamma, v)$. The image of p in Γ' is a path p' from v' to v' with the same label as that of p, that is $\mu(p) = \mu(p')$. However, p' need not be path-reduced. Namely, p' is path-reduced if and only if p does not contain any subpaths of the form e_2^{-1}, e_1 or e_1^{-1}, e_2. Let p'' be the path obtained from p' by performing all possible path reductions in Γ'. Then $\overline{\mu(p)} = \overline{\mu(p')} = \overline{\mu(p'')}$ and $\mu(p'') \in L(\Gamma, v')$. Thus we have shown that $\overline{L(\Gamma, v)} \subseteq \overline{L(\Gamma', v')}$.

Suppose now that p' is an arbitrary reduced path in Γ' from v' to v'. We claim that there is a reduced path in Γ from v to v with exactly the same label as that of p'. We will construct this path explicitly.

The occurrences of $e^{\pm 1}$ (if any) subdivide p' into a concatenation of the form:

$$p' = p_0 f_0 p_1 f_1 \dots f_k p_{k+1}$$

where $f_i = e^{\pm 1}$ and the paths p_i do not involve $e^{\pm 1}$.

Suppose that for some i we have $f_i = e$. Since p_i and p_{i+1} do not involve the edge e, they can also be considered as paths in Γ. Moreover, by the definition of folding, in the graph Γ the terminal vertex of p_i is joined with the initial vertex of p_{i+1} by either the edge e_1 or the edge e_2. We denote this edge by d_i (so that $d_i \in \{e_1, e_2\}$). Note that now $p_i d_i p_{i+1}$ is a reduced path in Γ with the same label as the path $p_i f_i p_{i+1}$ in Γ'.

Similarly, if for some i we have $f_i = e^{-1}$, we can find $d_i \in \{e_1^{-1}, e_2^{-1}\}$ such that $p_i d_i p_{i+1}$ is a reduced path in Γ with the same label as the path $p_i f_i p_{i+1}$ in Γ'.

Then

$$p = p_0 d_0 p_1 \dots d_k p_{k+1}$$

is a reduced path in Γ from v to v with the same label as p'. Thus $\mu(p') \in L(\Gamma, v)$, $\overline{\mu(p')} \in \overline{L(\Gamma, v)}$ and therefore $\overline{L(\Gamma', v')} \subseteq \overline{L(\Gamma, v)}$.

Hence $\overline{L(\Gamma', v')} \subseteq \overline{L(\Gamma, v)}$ and the lemma is proved. ☐

Before investigating further group-theoretic properties of labeled graph we need to introduce the following important geometric concept.

Definition 3.5 (Core graphs)**.** Let Γ be an X-digraph and let v be a vertex of Γ. Then the *core of* Γ *at* v is defined as:

$$Core(\Gamma, v) = \cup \{ p \mid \text{ where } p \text{ is a reduced path in } \Gamma \text{ from } v \text{ to } v \}$$

It is easy to see that $Core(\Gamma, v)$ is a connected subgraph of Γ containing v. If $Core(\Gamma, v) = \Gamma$ we say that Γ is a *core graph with respect to* v.

Example 3.6. The graph shown in Figure 3 is not a core graph with respect to v. However, it is a core graph with respect to the terminus of the edge labeled c. Both graphs shown in Figure 1 are core graphs with respect to any of their vertices.

The following lemma lists some obvious properties of core graphs.

Lemma 3.7. *Let $\Gamma' = Core(\Gamma, v)$. Then:*

1. *The subgraph Γ' of Γ is connected and contains the vertex v.*
2. *The graph Γ' has no degree-one vertices, except possibly for the vertex v.*
3. *The languages of Γ and Γ' at v coincide, that is $L(\Gamma, v) = L(\Gamma', v)$.*

Since languages of folded graphs consist of freely reduced words, Lemma 3.4 implies that for a folded X-digraph Γ the language $L(\Gamma, v)$ is a subgroup of $F(X)$. In fact, by the properties of core graphs, if $\Gamma' = Core(\Gamma, v)$ then $L(\Gamma, v) = L(\Gamma', v)$ is a subgroup of $F(X)$. We will now show that any finitely generated subgroup of $F(X)$ (and, later, even infinitely generated) can be obtained in such a way.

Proposition 3.8. *Let H be a finitely generated subgroup of $F(X)$. Then there exists a finite X-digraph Γ and a vertex v of Γ such that:*

1. *the graph Γ is folded and connected;*
2. *all vertices of Γ, except possibly for vertex v, have degree greater than one;*
3. *the degree of each vertex in Γ is at most $2\#(X)$;*
4. *the graph Γ is a core graph with respect to v.*
5. *the language of Γ is equal to H, that is $L(\Gamma, v) = H$.*

Proof. If $H = 1$ then the graph consisting of a single vertex obviously satisfies all the requirements of the proposition. Assume now that H is a nontrivial finitely generated subgroup of $F(X)$.

Let H be generated by the elements $h_1, \dots, h_m$ (where we think of each h_i as a freely reduced word in $\Sigma = X \cup X^{-1}$).

We define an X-digraph Γ_1 as follows. The graph Γ_1 is a wedge of m circles wedged at a vertex called v_1. The i-th circle is subdivided into $|h_i|$ edges which are oriented and labeled by X so that the label of the i-th circle (as read from v_1 to v_1) is precisely the word h_i.

Then any freely reduced word in $h_1, \dots, h_m$ is the label of a reduced path in Γ_1 from v to v. The converse is also obviously true. Thus $\overline{L(\Gamma_1, v_1)} = \langle h_1, \dots, h_m \rangle = H$. Note also that Γ_1 is connected by construction and has no vertices of degree one.

We define a sequence of graphs $\Gamma_1, \Gamma_2, \dots$ inductively as follows (see Figure 4 for a specific example). Suppose Γ_i is already constructed. It it is folded, terminate the sequence. Otherwise, let Γ_{i+1} be obtained from Γ_i by a folding.

Since Γ_1 is finite and a folding decreases the number of edges, this sequence terminates in finitely many steps with a folded graph Γ_k. Moreover, since a folding of a connected graph is connected, the graphs $\Gamma_1, \dots, \Gamma_k$ are connected. Let v_i be the image of v_1 in the graph Γ_i. By Lemma 3.4 we have $H = \overline{L(\Gamma_1, v_1)} = \overline{L(\Gamma_k, v_k)} = L(\Gamma_k, v_k)$, the last equality implied by the fact that Γ_k is folded.

Put $\Gamma = \Gamma_k$ and $v = v_k$. We have already shown that $L(\Gamma, v) = H$ and that Γ is a connected folded finite graph. Since Γ is folded, the degrees of its vertices are at most $2\#(X) = \#(X)$. We claim that Γ is in fact a core graph with respect to v. Indeed, suppose this is not so and there is a degree-one vertex u of Γ which is different from v. Let e be the unique edge of $\hat{\Gamma}$ with terminus u and let $x \in \Sigma$ be the label of e . There exists an edge e_1 of the graph Γ_1 such that the image of e_1 in $\Gamma = \Gamma_k$ is e (so that the label of e_1 is also x . Let u_1 be the terminus of e_1. Since $u \neq v = v_k$, we have $u_1 \neq v_1$. Recall the explicit construction of Γ_1 as the wedge of circles labeled $h_1, \dots, h_m$. It follows from this construction that there is a path p_1 in Γ_1 from v_1 to v_1 which passes through the edge e_1 and has freely reduced label w (namely, we can take $w = h_i^{\pm 1}$ for some i). Let p be the image of the path p_1 in Γ_k after performing all $k-1$ foldings. Then p is a path from v to v passing through e and with a freely reduced label w. However the vertex u (different from v) is of degree one. Therefore any path from v to v passing through e contains a subpath e, e^{-1} and hence cannot have a freely reduced label. This gives us a contradiction. Thus $Core(\Gamma, v) = \Gamma$ and the pair (Γ, v) satisfies all the requirements of the proposition. □

Lemma 3.9. *Let Γ be a connected folded X-digraph which is a core graph with respect to some vertex v. Let $H = L(\Gamma, v)$. Then:*

1. *For any initial segment w of a freely reduced word $h \in H$ there exists a unique reduced path p in Γ with origin v and label w.*
2. *For any reduced path p in Γ with origin v and label w, the word w is a subword of some freely reduced word $h \in H$.*

Proof. (1) The uniqueness of p follows from the fact that Γ is folded. The existence of p is also obvious. Indeed, we can take p to be the initial segment of length $|w|$ of the path in Γ from v to v with label h.

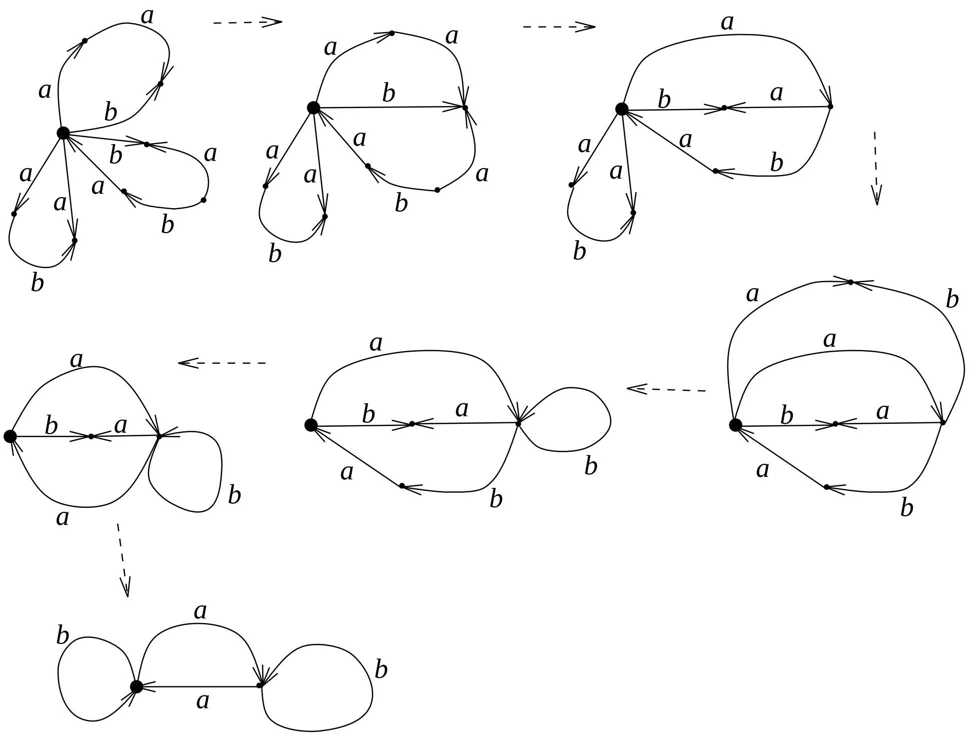

FIGURE 4. Constructing the graph associated to a subgroup. In this example we start with a subgroup $H = \langle a^2b, ba^{-1}ba, aba^{-1}\rangle \le F(a,b)$. The highlighted vertices represent the images of the original base-vertex v_1 in the intermediate graphs

(2) Let p be a path in Γ with origin v and label w. If the terminal vertex of p is v then $w \in H$ and the statement of the lemma is obvious. Suppose now that the terminal vertex u of p is different from v. Since Γ is a core graph with respect to v, the degree of u is at least two. Let e be an edge with origin u such that e^{-1} is different from the last edge in p. Let p' be a reduced path from $t(e)$ to v in Γ. Then $\alpha = pep'$ is a reduced path from v to v. Therefore $\mu(\alpha) = h \in H$. It is clear by construction that w is an initial segment of h. ☐

4. MORPHISMS OF LABELED GRAPHS

Recall that in Definition 2.1 we defined the notion of a morphism between two X-digraphs.

First we prove the following obvious but important lemma.

Lemma 4.1 (Morphisms and subgroups). *Let $\pi : \Gamma \longrightarrow \Gamma'$ be a morphism of X-digraphs such that $\pi(v) = v'$ for some vertex v of Γ. Suppose that Γ and Γ' are folded. Put $K = L(\Gamma, v)$ and $H = L(\Gamma', v')$. Then $K \le H$.*

Proof. Suppose p is an arbitrary reduced path in Γ from v to v and let w be the label of p. Note that w is freely reduced since Γ is folded and that $w \in L(\Gamma, v) = K$.

Since π is a morphism of X-digraphs, $\pi(p)$ is a path in Γ' from v' to v' with label w. Since w is freely reduced, the path $\pi(p)$ is path-reduced and therefore $w \in L(\Gamma', v') = H$. The path p was chosen arbitrarily and therefore $K \le H$. ☐

Lemma 4.2 (Uniqueness of morphisms for folded graphs)**.** *Let Γ and Δ be connected X-digraphs such that Δ is folded. Let v be a vertex of Γ and let u be a vertex of Δ. Then there exists at most one morphism of X-digraphs $f : \Gamma \longrightarrow \Delta$ such that $f(v) = u$.*

Proof. Suppose $f, j : \Gamma \longrightarrow \Delta$ are such that $f(v) = j(v) = u$. Let v' be an arbitrary vertex of Γ. Since Γ is connected, there exists a reduced path p in Γ from v to v'. Denote the label of p by w. Since Δ is folded, there exists at most one path in Δ with origin u and label w. On the other hand both $f(p)$ and $j(p)$ are such paths and they have terminal vertices $f(v')$ and $j(v')$ respectively. Therefore $f(v') = j(v')$. Since $v' \in V\Gamma$ was chosen arbitrarily, we have shown that f and j coincide on the vertex set of Γ. Since Γ is folded, it easily follows that f and j coincide on the set of edges as well. ☐

Proposition 4.3. *Let $F(X)$ be a free group with finite basis H. Let $K \le H \le F(X)$ be subgroups of $F(X)$. Suppose (Γ_1, v_1) and (Γ_2, v_2) are connected folded based X-digraphs such that Γ_i is a core graph with respect to v_i and $L(\Gamma_1, v_1) = K$ and $L(\Gamma_2, v_2) = H$.*

Then there exists a unique morphism of X-digraphs $\pi : \Gamma_1 \longrightarrow \Gamma_2$ such that $\pi(v_1) = v_2$.

Proof. The uniqueness of π follows from Lemma 4.2. Thus it suffices to show that such π exists.

We will construct π explicitly as follows. Let v be a arbitrary vertex of Γ_1. Choose any reduced path p_v in Γ from v_1 to v and let w be the label of p_v. By Lemma 3.9 since $L(\Gamma_1, v_1) = K$ and Γ_1 is a folded core graph with respect to v_1, the word w is an initial segment of some freely reduced word belonging to K. Since $K \le H$ and $H = L(\Gamma_2, v_2)$ Lemma 3.9 also implies that there is a unique path q_v in Γ_2 with origin v_2 and label w. We set $\pi(v)$ to be the terminal vertex of q_v. We claim that this definition of $\pi(v)$ does not depend on the choice of p_v. Indeed, let p'_v be another reduced path in Γ_1 from v_1 to v and let w' be the label of p'_v. Again, let q'_v be the path in Γ_2 with origin v_2 and label w'. Denote the terminal vertex of p_v by u and the terminal vertex of p'_v by u'. Note that $p_v(p'_v)^{-1}$ is a loop in Γ_1 at v_1 with label $w(w')^{-1}$. Hence $\overline{w(w')^{-1}} = k \in K \le H$. Let r be the reduced path in Γ_1 from v_1 to v_1 with label k. Then rp'_v is a path in Γ_1 from v_1 to u'. On the other hand the path-reduced form of rp'_v has label $k \cdot w' = w$. Since Γ_2 is folded, there is no more than one path in Γ_2 with origin v_1 and label w, namely the path p_v. Since path-reductions do not change end-points of paths, this means that $u' = v'$, as required.

We now extend π to the edges of Γ_1. Let e be an edge of Γ_1 with label x. Since Γ_2 is folded, there is at most one edge in Γ_2 with label x and origin $\pi(o(e))$. We claim that such an edge in fact exists. Since Γ_1 is a core graph with respect to 1_K, there is a reduced path $p = p'ep''$ from v_1 to v_1 passing through e. Let $z' = \mu(p'), z'' = \mu(p'')$ so that $z = \mu(p) = z'xz''$. By assumption on Γ_1 we have $z \in K$. However, $K \le H$, so that $z \in H$ as well. Hence there exists a unique path q is Γ_2 with origin v_2 and label $z = z'xz''$. This q has the form $q = y'e'y''$ where the label of q' is z', the label of edge e' is x and the label of y'' is z''. By our construction of π the terminus of y' is $\pi(o(e))$. Hence the edge e' has label x and origin $\pi(o(e))$ in Γ_2. We put $\pi(e) = e'$.

Thus we have established that π exists. ☐

Lemma 4.4 (Morphism factorization)**.** *Let $f : \Gamma \longrightarrow \Delta$ be an epimorphism of X-digraphs. Then f can be decomposed as $f = f_2 \circ f_1$ where $f_1 : \Gamma \longrightarrow \Gamma'$ is injective on the edge set of Γ and $f_2 : \Gamma' \longrightarrow \Delta$ is a bijection between the vertex sets of Γ' and Δ.*

Proof. Let $f : \Gamma \longrightarrow \Delta$ be an epimorphism of X-digraphs. Define a graph Γ' as follows. We give an informal description of Γ', f_1 andf_2 and leave the technicalities to the reader. The map f_1 consists in identifying those vertices (but not edges) of Γ which are identified by f. Thus $\Gamma' = f_1(\Gamma)$ is obtained by collapsing some subsets of $V\Gamma$ into single vertices.

Note now that if e, h are edges of Γ with the same image in Δ under f, then the origins of $f_1(e), f_1(h)$ are the same and the termini of $f_1(e), f_1(h)$ are the same. The morphism f_2 consists in identifying all such $f_1(e), f_1(h)$. Thus f_2 is in fact an identity map on the set of vertices of Γ'. Moreover, f_2 could be considered as a "generalized folding". In fact, f_2 is a composition of several foldings if Γ is finite. ☐

Convention 4.5. A *based X-digraph* is an X-digraph Γ with a marked vertex v, called the *base-vertex*. Such a based digraph is denoted (Γ, v). A *morphism* of based X-digraphs $\pi : (\Gamma, v) \longrightarrow (\Delta, u)$ is a morphism of the underlying X-digraphs $\pi : \Gamma \longrightarrow \Delta$ such that $\pi(v) = u$. By Lemma 4.2 if two folded connected based digraphs (Γ, v) and (Δ, u) are isomorphic then there exists a unique isomorphism $\pi : (\Gamma, v) \longrightarrow (\Delta, u)$. Therefore we can identify (Γ, v) and (Δ, u) via π. In this situation we will sometimes write $(\Gamma, v) = (\Delta, u)$.

5. Definition of subgroup graphs

In Proposition 3.8 we saw that one can associate to any finitely generated subgroup $H \leq F(X)$ a finite folded connected X-digraph Γ which is a core graph with respect to some vertex v and such that $L(\Gamma, v)$. In this section we show that the same can be done for an arbitrary subgroup $H \leq F(X)$ (although the graph will not be finite if H is not finitely generated) and the pair (Γ, v) is unique for a fixed H.

Theorem 5.1. *Let $F(X)$ be a free group with finite basis X. Let $H \leq F(X)$ be a subgroup of $F(X)$.*

Then there exist a connected folded X-digraph Γ and a vertex v of Γ such that:

1. *the graph Γ is a core graph with respect to v;*
2. *$L(\Gamma, v) = H$.*

Proof. First we will construct an X-digraph Δ as follows.

The vertex set of Δ is the set of left cosets of H in $F(X)$:

$$V\Delta = \{Hf \mid f \in F(X)\}.$$

For two cosets Hf and Hg and a letter $x \in X$ we introduce a directed edge with origin Hf, terminus Hg and label x whenever $Hfx = Hg$. This defines an X-digraph Δ. Put $v = H$ to be the coset of the identity element $1 \in F(X)$.

Note that Δ is a connected graph. Indeed, suppose $f = x_1 \dots x_n$ is an arbitrary nontrivial freely reduced word in $X^{\pm 1}$ (here $x_i \in X \cup X^{-1}$). Then for each $i = 1, \dots, n-1$ there is an edge in $\hat{\Delta}$ from $Hx_1 \dots x_i$ to $Hx_1 \dots x_i x_{i+1}$ with label x_{i+1}. Therefore there is a path in $\hat{\Delta}$ from $v = H1$ to Hf with label f, and so Δ is connected.

Observe now that Γ is folded. Suppose this is not the case and there is a vertex $u = Hf$ with two distinct edges e_1, e_2 which have the same label $x \in X \cup X^{-1}$. Let u_1 and u_2 be the terminal vertices of e_1 and e_2 accordingly. By definition of Δ this means that $u_1 = fxH = u_2$. However, by the construction of Δ, for any pair of vertices v_1, v_2 and any $x \in X \cup X^{-1}$ there is at most one edge from v_1 to v_2 with label x. Thus $e_1 = e_2$, contrary to our assumptions.

We claim that $L(\Delta, v) = H$. Indeed, suppose f is a nontrivial freely reduced word with is a label of a loop p at v in Δ. Again let $f = x_1 \dots, x_n$ where $x_i \in X \cup X^{-1}$. The definition of Δ implies that for every $i = 1, \dots, n$ the terminal vertex of the initial segment of p with label $x_1 \dots x - i$ is the coset $Hx_1 \dots x_i$. In particular the terminal vertex of p is $Hx_1 \dots x_n = Hf$. On the other hand the terminal vertex of p is H by assumption. Therefore $Hf = H$ and so $f \in H$. Thus we have shown that $L(\Delta, v) \subseteq H$.

Suppose now that f is an arbitrary nontrivial freely reduced word such that $f \in H$. As we have seen when proving the connectivity of Δ, there exists a path p in Δ with origin $v = H$ and label f. Obviously p is reduced since f is freely reduced and Δ is folded. Let u be the terminal vertex of p. By the construction of p we have $u = Hf = H = v$, since $f \in H$. Thus p is a reduced loop at v with label f. Hence $f \in L(\Delta, v)$. Since $f \in H$ was chosen arbitrarily, we have proved that $H \subseteq L(\Delta, v)$.

Thus $H \subseteq L(\Delta, v) \subseteq H$ and therefore $H = L(\Delta, v)$.

Now put $\Gamma = Core(\Delta, v)$. It is obvious that Γ satisfies all the requirements of Theorem 5.1. $\square$

In a seminal paper [42] C.Sims uses a similar approach to that used in the proof of Theorem 5.1 to study subgroups of free groups via "important" parts of their coset graphs.

Theorem 5.2. *Let $F(X)$ be a free group with finite basis X and let $H \leq F(X)$ be a subgroup of $F(X)$. Suppose (Γ_1, v_1) and (Γ_2, v_2) are connected folded based X-digraphs such that Γ_i is a core graph with respect to v_i and $L(\Gamma_1, v_1) = H = L(\Gamma_2, v_2)$.*

Then there exists a unique isomorphism of X-digraphs $\pi : \Gamma_1 \longrightarrow \Gamma_2$ such that $\pi(v_1) = v_2$.

Proof. The uniqueness of π follows from Lemma 4.2.

Since $H \leq H$, by Proposition 4.3 there is a morphism $\pi : \Gamma_1 \longrightarrow \Gamma_2$ such that $\pi(v_1) = v_2$. We claim that π is an isomorphism of X-digraphs. Indeed, Proposition 4.3 implies that there is a morphism $\pi' : \Gamma_2 \longrightarrow \Gamma_1$ such that $\pi'(v_2) = v_1$. Therefore $\pi' \circ \pi : \Gamma_1 \longrightarrow \Gamma_1$ is a morphism such that $(\pi' \circ \pi)(v_1) = v_1$. By Lemma 4.2 there is at most one such morphism, namely the identity map on Γ_1. Thus $\pi' \circ \pi = Id_{\Gamma_1}$. By a symmetric argument $\pi \circ \pi' = Id_{\Gamma_2}$. Therefore π is an isomorphism as required. □

Definition 5.3 (Subgroup graph)**.** Let $H \leq F(X)$ be a subgroup of $F(X)$. Then by Theorem 5.1 and Theorem 5.2 there exists a unique (up to a canonical isomorphism of based X-digraphs) based X-digraph (Γ, v) such that

1. the graph Γ is folded and connected;
2. the graph Γ is a core graph with respect to v;
3. the language of Γ with respect to v is H, that is $L(\Gamma, v) = H$.

In this situation we call Γ the *subgroup graph* (or a *subgroup graph-automaton*) of H with respect to X and denote it by $\Gamma(H)$ (or $\Gamma^X(H)$). The base-vertex v is denoted 1_H.

Lemma 5.4. *Let $H \leq F(X)$ be a subgroup of $F(X)$. Then the graph $\Gamma(H)$ is finite if and only if H is finitely generated.*

Proof. If H is finitely generated then $\Gamma(H)$ is finite by Proposition 3.8 and Proposition 4.3. If $\Gamma(H)$ is finite then H is finitely generated by Lemma 6.1. □

6. Spanning trees and free bases of subgroups of free groups

Lemma 3.3 provides a generating set for the subgroup corresponding to a X-digraph. However, one would like to be able to find a free basis or even a free Nielsen reduced basis for this subgroup. We will show next that both these goals can be easily accomplished using the original X-digraph.

Recall that in a connected graph a subgraph is said to be *a spanning tree* if this subgraph is a tree and it contains all vertices of the original graph. If a graph T is a tree then for any two vertices u, u' of T there is a unique reduced path in T from u to u' which will be denoted by $[u, u']_T$.

Lemma 6.1. *Let Γ be a folded X-digraph and let v be a vertex of Γ. Let T be a spanning tree of Γ. Let T^+ be the set of those positive edges of Γ which lie outside of T. For each $e \in T^+$ put $p_e = [v, o(e)]_T e[t(e), v]$ (so that p_e is a reduced path from v to v and its label is a freely reduced word in $\Sigma = X \cup X^{-1}$.) Also for each $e \in T^+$ put $[e] = \mu(p_e) = \overline{\mu(p_e)}$. Denote*

$$Y_T = \{[e] \mid e \in T^+\}.$$

Then Y_T is a free basis for the subgroup $H = L(\Gamma, v)$ of $F(X)$.

Proof. We can extend the definition of p_e and $[e]$ for all edges of $\hat{\Gamma}$ (whether positive or negative and whether inside or outside of T) by putting $p_e = [v, o(e)]_T e[t(e), v]$ and $[e] = \overline{\mu(p_e)}$. Note that in this case $p_{e^{-1}} = (p_e)^{-1}$ and $[e^{-1}] = [e]^{-1}$.

It is easy to see that if $e \in E^+T$ then $p_e = [v, o(e)]_T e[t(e), v]$ is a path that can be transformed by path-reductions into the trivial path and so $[e] = 1$. We know from Lemma 3.3 that the subgroup H is generated by the set $\{[e] \mid e \in E^+(\Gamma)\}$. Therefore H is in fact generated by the set $Y_T = \{[e] \mid e \in E^+T\}$. It remains to show that Y_T is in fact a free basis for H.

To see this it suffices to show that any nontrivial freely reduced word in $Y_T^{\pm 1}$ defines a nontrivial element of $F(X)$. Suppose $h = [e_1] \cdot \dots \cdot [e_k]$ where $k \geq 1$, $e_i \in Y^+ \cup (Y^+)^{-1} = E(\hat{\Gamma} - T)$ and $e_i \neq e_{i+1}^{-1}$. We need to show that $h \neq 1$. Since T is a tree, for any vertices u, u' of Γ the path-reduced form of the path $[u, v]_T[v, u']_T$ is the path $[u, u']_T$. By definition of h and of $[e_i]$ we have $h = \overline{p}$ where p is the following path from v to v in Γ:

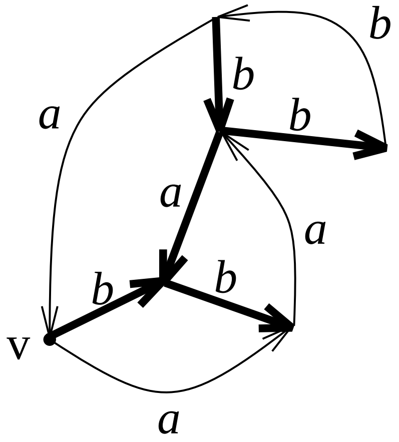

FIGURE 5. Free basis corresponding to a spanning tree. In this figure the spanning tree T is highlighted. The free basis Y_T for the subgroup $H = L(\Gamma, v)$ is $\{a^{-1}bab^{-1}, ba^{-1}b^3ab^{-1}, b^2a^2b^{-1}, ab^{-2}\}$.

$$p = [v, o(e_1)]_T e_1 [t(e_1), v]_T [v, o(e_2)]_T e_2 \dots [v, o(e_k)]_T e_k [t(e_k), v].$$

Also by the remark above the path p can be transformed by path-reductions to the path

$$p' = [v, o(e_1)]_T e_1 [t(e_1), o(e_2)]_T e_2 \dots e_i [t(e_i), o(e_{i+1})]_T e_{i+1} \dots e_k [t(e_k), v]_T.$$

We claim that the path p' is already path-reduced. Note that the segments $[t(e_i), o(e_{i+1})]_T$, $[v, o(e_1)]_T$, $[t(e_k), v]_T$ are contained in the tree T and the edges e_i are outside of T. Hence if p' is not path-reduced, then for some i $t(e_i) = o(e_{i+1})$ and $e_i = e_{i+1}^{-1}$. However this is impossible by our assumption on h.

Thus p' is indeed a nontrivial reduced path in Γ. Since Γ is folded, this implies that the label of p' is a nontrivial freely reduced word in X which therefore represents a nontrivial element of $F(X)$. But p' was obtained by a series of path-reductions from p and hence

$$1 \neq \mu(p') = \overline{\mu(p')} = \overline{\mu(p)} = h$$

and $h \neq 1$ as required. □

Proposition 6.2. *Let Γ be a folded connected X-digraph and let Γ' be a connected subgraph of Γ. Let v be a vertex of Γ'.*

Then $H = L(\Gamma', v)$ is a free factor of $G = L(\Gamma, v)$. Further, if Γ' does not contain $Core(\Gamma, v)$, then $H \neq G$. In particular, if both Γ and Γ' are core graphs with respect to v and $\Gamma' \subsetneqq \Gamma$ then $H \neq G$.

Proof. Let T' be a spanning tree in Γ'. Then there exists a spanning tree T of Γ such that T' is a subgraph of T. By Lemma 6.1 the set $Y_T = \{[e] \mid e \in E^+(\Gamma - T)\}$ is a free basis of G. Moreover, Y_T is a disjoint union of the sets:

$$\begin{aligned} Y_T =& \{[e] \mid e \in E^+(\Gamma - T), e \in \Gamma'\} \cup \{[e] \mid e \in E^+(\Gamma - T), e \notin E\Gamma'\} = \\ =& \{[e] \mid e \in E^+(\Gamma' - T')\} \cup \{[e] \mid e \in E^+(\Gamma - T), e \notin E\Gamma'\} = Y_{T'} \cup Z \end{aligned}$$

Thus $Y_{T'}$ is a subset of Y_T. Since Y_T is a free basis of G and $Y_{T'}$ is a free basis of G, this implies that

$$G = F(Y_T) = F(Y_{T'}) * F(Z) = H * F(Z)$$

and H is a free factor of G as required.

Suppose now that Γ' does not contain $Core(\Gamma, v)$. We claim that there is a positive edge e of Γ which does not belong to Γ' and is not in T. Assume this is not the case. Then all edges outside of Γ' lie in T. Hence $\Gamma - \Gamma' \subset T$ is a union of disjoint trees. This implies that $Core(\Gamma, v)$ is contained in Γ', contrary to our assumptions. Thus the claim holds and hence $Z \neq \emptyset$. Therefore $H \neq G$ as required. □

Corollary 6.3. *Let Γ be a connected folded X-digraph which is a core graph with respect to a vertex v. Let Γ' be a connected subgraph of Γ, containing v. Suppose that Γ' is also a core graph with respect to v. Then $L(\Gamma', v)$ is a free factor of $L(\Gamma, v)$ and $L(\Gamma', v) \neq L(\Gamma, v)$.*

We will show next that the procedure described in Lemma 6.1 can be improved to obtain a Nielsen-reduced free basis of a finitely generated subgroup of $F(X)$. Here we need to recall the following important definition.

Definition 6.4 (Nielsen set)**.** Let S be a set of nontrivial elements of the free group $F(X)$ such that $S \cap S^{-1} = \emptyset$. We say that S is *Nielsen reduced with respect to the free basis X* if the following conditions hold:
(1) If $u, v \in S \cup S^{-1}$ and $u \neq v^{-1}$ then $|u \cdot v|_X \geq |u|_X$ and $|u \cdot v|_X \geq |v|_X$.
(2) If $u, v, w \in S \cup S^{-1}$ and $u \neq w^{-1}, v \neq w^{-1}$ then $|u \cdot w \cdot v|_X > |u|_X + |v|_X - |w|_X$.

Condition (1) means that no more than a half of u and no more than a half of v freely cancels in the product $u \cdot v$. Condition (2) means that at least one letter of w survives after all free cancellations in the product $u \cdot w \cdot v$.

Recall that the vertex set of any connected graph Γ is canonically endowed with an integer-valued metric d_Γ. Namely, the distance between any two vertices u, u' is defined as the smallest length of an edge-path from u to v in Γ. An edge-path whose length is equal to the distance between its endpoints is said to be *geodesic in* Γ.

Definition 6.5 (Geodesic tree)**.** Let Γ be a connected graph with a base-vertex v. A subtree tree T in Γ is said to be *geodesic relative to* v if $v \in T$ and for any vertex u of T the path $[v, u]_T$ is geodesic in Γ, that is a path of the smallest possible length in Γ from v to u.

It is easy to see that geodesic spanning trees always exist:

Lemma 6.6. *Let Γ be a graph (whether finite or infinite) with a base-vertex v. Then there exists a geodesic relative to v spanning tree T for Γ.*

Proof. We will construct the tree T inductively. If Γ is finite or locally finite and recursive, our procedure will produce an actual algorithm for building T.

We will construct by induction a (possibly finite) sequence of nested trees $T_0 \subseteq T_1 \subseteq \dots$ in Γ such that for each $n \geq 0$
(1) the vertex set of T_n is precisely the ball $B(v, n)$ of radius n centered at v in the metric space $(V\Gamma, d_\Gamma)$ and
(2) the subtree T_n is geodesic relative to v in Γ.
Step 0**.** Put $T_0 = v$. Obviously T_0 is a tree and $T_0 = B(v, 0)$.
Step n**.** Suppose that the trees $T_0 \subseteq T_1 \subseteq \dots \subseteq T_{n-1}$ have already been constructed. Suppose further that for each $i \leq n - 1$ we have $VT_i = B(v, i)$ and that T_i is geodesic relative to v.

For each vertex u of Γ at the distance n from v choose an edge e_u with terminus u such that $d(o(e_u), v) = n - 1$. (For instance we can take e_u to be the last edge of a geodesic path from v to u in Γ). Put

$$T_n = T_{n-1} \cup \bigcup \{e_u \,|\, u \in V\Gamma, d(v, u) = n\}$$

It is easy to see that T_n is a geodesic tree containing T_{n-1} and that $VT_n = B(v, n)$.

Put $T = \underset{n \geq 0}{\cup} T_n$. Then T is a geodesic spanning tree (relative to v) for Γ, as required. □

Proposition 6.7 (Nielsen basis)**.** *Let Γ be a folded X-digraph which is a core graph with respect to a vertex v of Γ. Let $H = L(\Gamma, v) \le F(X)$ and let T be a spanning tree in Γ which is geodesic with respect to v.*

Then the set Y_T is a Nielsen-reduced free basis of the subgroup H.

Proof. Recall that $Y_T = \{[e] = \mu([v, o(e)]_T e[t(e), v]_T) \,|\, e \in E^+(\Gamma - T)\}$. Moreover, $[e]^{-1} = [e^{-1}] = \mu([v, o(e^{-1})]_T e[t(e^{-1}), v]_T)$ where $e \in E^+(\Gamma - T)\}$.

Note that since the tree T is geodesic we have $||[v, o(e)]_T| - |[t(e), v]_T|| \le 1$ for each $e \in E(\Gamma - T)$. Note also that the path $p_e = [v, o(e)]_T e[t(e), v]_T$ (where $e \in E(\Gamma - T)$) is path-reduced and therefore its label is freely reduced.

This means that for the path p_e (with $e \in E(\Gamma - T)$) we have

$$|[v, o(e)]_T| \le (1/2)|p_e| = (1/2)|[e]|_X \quad \text{and} \quad |[t(e), v]_T| \le (1/2)|p_e| = (1/2)|[e]|_X$$

and

$$|e[t(e), v]_T| \ge (1/2)|p_e| = (1/2)|[e]|_X, \quad |[v, o(e)]_T e| \ge 1/2)|p_e| = (1/2)|[e]|_X.$$

We already know from Lemma 6.1 that the set Y_T is a free basis of H. It remains to check that Y_T is Nielsen-reduced.

We first check condition (1) of Definition 6.4. Suppose $e, f \in E(\Gamma - T)$ and $e \ne f^{-1}$. Then the path-reduced form of the path $p_e p_f = [v, o(e)]_T e[t(e), v]_T [v, o(f)]_T f[t(f), v]_T$ is the path

$$p' = [v, o(e)]_T e[t(e), o(f)]_T f[t(f), v]_T.$$

The label of this path is freely reduced (since Γ is folded and therefore it is equal to $[e] \cdot [f] \in H$. Since p' has a subpath $[v, o(e)]_T e$, we conclude that $|[e] \cdot [f]|_X \ge (1/2)|[e]|_X$. Since p' has a subpath $f[t(f), v]_T$, we conclude that $|[e] \cdot [f]|_X \ge (1/2)|[f]|_X$. Thus condition (1) holds.

We will now verify condition (2) of Definition. Suppose $e, f, g \in E(\Gamma - T)$ are such that $e \ne f^{-1}, g \ne f^{-1}$. We need to show that $|[e] \cdot [f] \cdot [g]|_X > |[e]|_X + |[g]|_X - |[f]|_X$.

The path-reduced form of the concatenation

$$p_e p_f p_g = [v, o(e)]_T e[t(e), v]_T [v, o(f)]_T f[t(f), v]_T [v, o(g)]_T g[t(g), v]_T$$

is the path

$$p'' = [v, o(e)]_T e[t(e), o(f)]_T f[t(f), o(g)]_T g[t(g), v]_T$$

When $p_e p_f p_g$ is transformed to p', we have to cancel the terminal segment of p_e of length at most $|[v, o(f)]_T|$ and the initial segment of p_g of length at most $|[t(f), v]_T|$. Since the edge f of p_f survives, we have

$$|p''| \ge 1 + |p_e| + |p_g| - |[v, o(f)]_T| - |[t(f), v]_T| \ge 1 + |p_e| + |p_g| - (1/2)|p_f| - (1/2)|p_f| > |p_e| + |p_g| - |p_f|$$

and so

$$|[e] \cdot [f] \cdot [g]|_X > |[e]|_X + |[g]|_X - |[f]|_X$$

as required. □

Corollary 6.8. *Let $H \le F(X)$ be a subgroup of $F(X)$. Then H is free. Moreover, there is a Nielsen-reduced free basis for H.*

Proof. Consider the graph $\Gamma(H)$. By Lemma 6.6 there exists a geodesic spanning tree T in $\Gamma(H)$. Hence by Proposition 6.7 the set Y_T is a Nielsen-reduced free basis for H. In particular H is free. □

Historical Note 6.9. The fact that finitely generated subgroups in a free group $F(X)$ of finite rank are free must have been known to Klein and Fricke in 1880's. The first formal proof is due to Jacob Nielsen and appeared in 1923 (see [38]). Later Nielsen generalized it to arbitrary subgroups of $F(X)$ in 1955 [36]. Today this statement appears trivial by topological considerations. Indeed, $F(X)$ can be realized as the fundamental group of a topological graph. A covering space of a graph is a graph and fundamental groups of graphs are free.

However, the existence of a *Nielsen-reduced* free basis is not at all trivial, even for finitely generated subgroups of $F(X)$. The standard argument, due to J.Nielsen himself, uses a rather complicated machinery of Nielsen transformations and a particular choice of well-ordering on $F(X)$. The advantage of our approach is that we use only very elementary combinatorial objects.

Remark 6.10. It can be shown, for example by looking at abelianizations, that for a free group any two bases have the same cardinality. If F is a finitely generated free group, then the cardinality of any free basis of F is called *the rank of* F and denoted $rk(F)$. It is easy to see that for finitely generated free groups F and F_1 the free product $F * F_1$ is also a free group and $rk(F * F_1) = rk(F) + rk(F_1)$.

7. Basic algebraic and algorithmic properties of subgroup graphs

In this section we will restrict our attention to finitely generated subgroups of free groups. Moreover, we will be particularly interested in algorithmic aspects of various group-theoretic statements.

Till the end of this section let X be a finite set and let $F = F(X)$ be the free group on X.

Proposition 7.1 (Constructing $\Gamma(H)$)**.** *There is an algorithm which, given finitely many freely reduced X-words $h_1, \dots, h_k$, constructs the graph $\Gamma(H)$, where $H = \langle h_1, \dots, h_k \rangle$.*

Proof. This follows immediately from the proof of Proposition 3.8. Indeed, if we start with a wedge of k circles with the words $h_1, \dots, h_k$ written on them and then perform all possible foldings, the resulting X-digraph will be $\Gamma(H)$. Note that the original wedge-of-circles graph has $M = |h_1|_X + \cdots + |h_k|_X$ edges. Since every folding reduces the number of edges by one, we will obtain the graph $\Gamma(H)$ in at most M steps. ☐

Proposition 7.2 (Generalized word problem)**.** *The group $F = F(X)$ has* solvable generalized word problem*.*

That is, there exists an algorithm which, given finitely many freely reduced X-words $h_1, \dots, h_k$ and given a freely reduced X-word g decides whether or not g belongs to the subgroup $H = \langle h_1, \dots, h_n \rangle$ of F.

Proof. We first construct the graph $\Gamma(H)$ as it is done in Proposition 7.1. Recall that $L(\Gamma(H), 1_H) = H$ by definition of $\Gamma(H)$.

Thus to decide if $g \in H$, we simply have to check whether or not the graph $\Gamma(H)$ accepts the word g. Recall that $\Gamma(H)$ is folded and that the language $L(\Gamma(H), 1_H)$ consists of all labels of reduced paths in $\Gamma(H)$ from 1_H to 1_H. Given a freely reduced word $g = x_1 \dots x_m$, $x_i \in \Sigma = X \cup X^{-1}$ we can check whether $g \in L(\Gamma(H), 1_H)$ as follows.

First we check if there is an edge in $\hat{\Gamma}(H)$ with origin 1_H and label x_1 (since Γ is folded, such an edge is unique if it exists). If yes, we move to the terminal vertex of this edge, which we denote v_1. If no, we terminate the process. We then check if there is an edge in $\hat{\Gamma}(H)$ with origin v_1 and label x_2. If yes, we move to the terminal vertex of this edge, which we denote v_2. If no, we terminate the process. By repeating this procedure at most k times we either will terminate the process, in which case we conclude that $g \notin L(\Gamma(H), 1_H)$, or we find a vertex v_k such that the word g is the label of a reduced path in $\Gamma(H)$ from 1_H to v_k. If $v_k \neq 1_H$, we conclude that $g \notin L(\Gamma(H), 1_H)$ and if $v_k = 1_H$, we conclude that $g \in L(\Gamma(H), 1_H)$. ☐

Recall that a core graph is connected and has at most one vertex of degree one (namely the base-vertex).

Definition 7.3 (Type of a core graph)**.** Let Γ be a folded X-digraph which is a core graph with respect to some vertex. Suppose that Γ has at least one edge.

If every vertex of Γ has degree at least two, we set the *type* of Γ, denoted $Type(\Gamma)$ to be equal to Γ.

Suppose now that Γ has a vertex v of degree one (such a vertex is unique since Γ is a core graph). Then there exists a unique vertex v' of Γ with the following properties:

(a) There is a unique Γ-geodesic path $[v, v']$ from v to v' and every vertex of this geodesic, other than v and v', has degree two.
(b) The vertex v' has degree at least three.

Let Γ' be the graph obtained by removing from Γ all the edges of $[v, v']$ and all the vertices of $[v, v']$ except for v'. Then Γ' is called the *type of* Γ and denoted $Type(\Gamma)$.

Finally, if Γ consists of a single vertex, we set $Type(\Gamma) = \Gamma$.

Remark 7.4. Observe that $Type(\Gamma)$ is an X-digraph which does not have a distinguished base-vertex. Note also that the graph $\Gamma' = Type(\Gamma)$ is a core-graph with respect to any of its vertices. Moreover in the graph Γ' every vertex has degree at least two and $\Gamma' = Type(\Gamma')$.

Lemma 7.5. *Let Γ be a folded core graph (with respect to one of its vertices). Let v and u be two vertices of Γ and let q be a reduced path in Γ from v to u with label $g \in F(X)$. Let $H = L(\Gamma, v)$ and $K = L(\Gamma, u)$.*

Then $H = gKg^{-1}$.

Proof. Let p be a reduced path in Γ from u to u (so that $\mu(p) = k \in K$). Then the path $p' = qpq^{-1}$ is a path from v to v with label $\mu(p') = \mu(q)\mu(p)(\mu(q))^{-1}$. Thus the freely reduced form of $\mu(p')$ is equal to the element $g \cdot k \cdot g^{-1} \in F(X)$. The path p' can be transformed by several path-reductions to a reduced path p'' from v to v. Hence $\mu(p'') \in L(\Gamma, v) = H$. On the other hand $\mu(p'') = \overline{\mu(p')} = g \cdot k \cdot g^{-1}$. Thus we have shown that for any $k \in K$ we have $g \cdot k \cdot g^{-1} \in H$ so that $gKg^{-1} \subseteq H$. A symmetric argument shows that $g^{-1}Hg \subseteq K$ and thus $H \subseteq gKg^{-1}$. Therefore $H = gKg^{-1}$, as required.

□

Lemma 7.6. *Let $H \leq F(X)$ and let $\Gamma = \Gamma(H)$.*

Let $g \in F(X)$ be a nontrivial freely reduced word in X. Let $g = yz$ where z is the maximal terminal segment of the word g such that there is a path with label z^{-1} in Γ starting at 1_H (such a path is unique since Γ is folded). Denote the end-vertex of this path by u. Let Δ' be the graph obtained from Γ as follows. We attach to Γ at u the segment consisting of $|y|$ edges with label y^{-1}, as read from u. Let u' be the other end of this segment. Put $\Delta'' = Core(\Gamma', u')$.

Then $(\Delta'', u') = (\Gamma(K), 1_K)$ where $K = gHg^{-1}$.

Proof. Let Δ be the graph obtained by attaching to Γ at 1_H a segment of $|g|_X$ edges, labeled g, which has origin v and terminus 1_H. It is easy to see that $\overline{(}L(\Delta, v)) = gHg^{-1}$. However, the graph Δ is not necessarily folded. If we fold the terminal segment z of g (as defined in the statement of Lemma 7.6) onto the path in Γ with origin 1_H and label z^{-1}, the resulting graph Δ' is obviously folded. Therefore $L(\Delta', u') = H$ (where u' is the image of v in Δ'). It can still happen that Δ' is not a core graph with respect to u'. (see Figure). However the graph $\Delta'' = Core(\Delta', u')$ is a folded, is a core graph with respect to u' and has the property $L(\Delta'', u') = K$. Hence $(\Delta'', u') = (\Gamma(K), 1_K)$, as required. □

The following observation will allow us to decide when two finitely generated subgroups of $F(X)$ are conjugate.

Proposition 7.7. *[Conjugate subgroups] Let H and K be subgroups of $F(X)$.*

Then H is conjugate to K in $F(X)$ if and only if the graphs $Type(\Gamma(H))$ and $Type(\Gamma(K))$ are isomorphic as X-digraphs.

Proof. Suppose that $Type(\Gamma(H)) = Type(\Gamma(K)) = \Gamma$. Let v be a vertex of Γ. It follows from Lemma 7.5 that the subgroup $L(\Gamma, v)$ is conjugate to both H and K, so that H is conjugate to K.

Suppose now that K is conjugate to H, that is $K = gHg^{-1}$ for some $g \in F(X)$. Lemma 7.6 implies that $Type(\Gamma(H)) = Type(\Gamma(K))$, as required. □

Corollary 7.8 (Conjugacy problem for subgroups of free groups)**.** *There is an algorithm which, given finitely many freely reduced words $h_1, \dots, h_s, k_1, \dots, k_m$, decides whether or not the subgroups $H = \langle h_1, \dots, h_s \rangle$ and $K = \langle k_1, \dots, k_m \rangle$ are conjugate in $F(X)$*

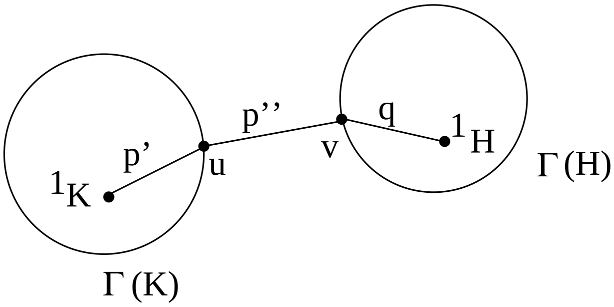

FIGURE 6

Proof. First we construct the graphs $\Gamma(H)$ and $\Gamma(K)$, as described in Proposition 7.1. By chopping-off segments with origin 1_H, 1_K if necessary, we then obtain the graphs $\Gamma = \mathit{Type}(\Gamma(H))$ and $\Delta = \mathit{Type}(\Gamma(K))$. It remains to check whether Γ and Δ are isomorphic as X-digraphs. This can be done by fixing a vertex $v \in V\Gamma$ and comparing for each vertex $u \in V\Delta$ the based X-digraphs (Γ, v) and (Δ, u). ☐

Proposition 7.9. *Let $H, K \leq F(X)$ and $g \in F(X)$. Let $g = yz$ where z is the maximal terminal segment of g such that z^{-1} is the label of a path in $\Gamma(H)$ with origin 1_H. Let $y = y'y''$ where y' is the maximal initial segment of y which is the label of a path in $\Gamma(K)$ with origin 1_K.*

Suppose that the word y'' is nontrivial. Then

$$\langle gHg^{-1}, K\rangle = gHg^{-1} * K.$$

Proof. Consider the graph Γ obtained in the following way. First we take an edge-path p_g of length $|g|_X$ with label g and attach its origin to 1_K and its terminus to 1_H. The path p_g is a concatenation of the form $p_g = p'p''q$ where $\mu(p') = y'$, $\mu(p'') = y''$ and $\mu(q) = z$.

Then we fold the segment q of p_g onto the path (starting at 1_H) with label z^{-1} in $\Gamma(H)$ and we fold the segment y' of p_g onto the path (starting at 1_K) in $\Gamma(K)$. We denote the resulting graph Γ.

Thus Γ consists of the graphs $\Gamma(H)$ and $\Gamma(K)$ joined by a segment with label y'' (see Figure 6). By a slight abuse of notation we will still call this segment p'' and we will also refer to the images of $\Gamma(H)$ and $\Gamma(K)$ in Γ as $\Gamma(H)$ and $\Gamma(K)$. The construction of Γ implies that it is a folded X-digraph.

It is not hard to see now that $L(\Gamma, 1_K) = gHg^{-1} * K$. Indeed, choose a spanning tree T_H is $\Gamma(H)$ and a spanning tree T_K in $\Gamma(K)$. Then $T = T_H \cup T_K \cup p''$ is a spanning tree in Γ. We can use it to produce a free basis Y_T for $L(\Gamma, 1_K)$ as described in Lemma 6.1. Then Y_T is a disjoint union of two subsets $Y_T = Y_H \sqcup Y_K$ where Y_H are the generators corresponding to the edges in $\Gamma(H) - T_H$ and Y_K are the generators corresponding to the edges in $\Gamma(K) - T_K$.

Each $f \in Y_H$ is the label of a path $\alpha_f = [1_K, u]_{T_K} p''[v, o(e)]_{T(H)} e[t(e), v]_{T_H} (p'')^{-1}[u, 1_K]_{T_K}$ where $e \in E^+(\Gamma(H) - T_H)$. The path p_f can be obtained by path-reductions from

$$\alpha'_f = ([1_K, u]_{T_K}(p')^{-1})p'p''q(q^{-1}[v, 1_H]_{T_H})[1_H, o(e)]_{T_H} e[t(e), 1_H]_T([1_H, v]_{T_H} q)q^{-1}(p'')^{-1}(p')^{-1}(p'[u, 1_K]_{T_K})$$

that is

$$\alpha'_f = ([1_K, u]_{T_K}(p')^{-1})p_g(q^{-1}[v, 1_H]_{T_H})p_e([1_H, v]_{T_H} q)p_g^{-1}(p'[u, 1_K]_{T_K})$$

where $p_e = [1_H, o(e)]_{T_H} e[t(e), 1_H]_T$ is the path from 1_H to 1_H in $\Gamma(H)$ defining one of the standard free generators of H relative the tree T_H.

Since $[1_K, u]_{T_K}(p')^{-1}$ is a path in $\Gamma(K)$ from 1_K to 1_K, the label of this path determines an element $k \in K$. Similarly, since $q^{-1}[v, 1_H]_{T_H} = ([1_H, v]_{T_H} q)^{-1}$ is a path in $\Gamma(H)$ from 1_H to 1_H, the freely reduced form of its label is an element $h \in H$.

Since the set $Y_{T_H} = \{\mu(p_e) \mid e \in E^+(\Gamma(H) - T_H)\}$ is a free basis of H, the set $Z_H = g \cdot h \cdot Y_{T_H} \cdot h^{-1} \cdot g^{-1}$ is a free basis of gHg^{-1}. Thus $Y_H = \overline{k \cdot Z_H k^{-1}}$.

On the other hand the set Y_K is clearly exactly the free basis of K corresponding to the spanning tree T_K in $\Gamma(K)$.

Since we know that $Y_T = Y_H \cup Y_K$ is a free basis of the subgroup $L(\Gamma, 1_K)$ and $F(Y_K) = K$ it follows that $(k^{-1}Y_H k) \cup Y_T = Z_H \cup Y_T$ is also a free basis of $L(\Gamma, 1_K)$. Therefore

$$L(\Gamma, 1_K) = F(Z_H \cup Y_K) = F(Z_H) * F(Y_K) = gHg^{-1} * K$$

as required. □

Proposition 7.10. *Let $H \leq F(X)$ be a finitely generated subgroup and let $\Gamma = \Gamma(H)$. Let $g \in F(X)$ be such that $g^n \in H$ for some $n \geq 1$. Then there exists $m \leq \# V\Gamma$, $m > 0$, such that $g^m \in H$.*

Proof. We may assume that $g \notin H$, since otherwise the statement is obvious.

Let $g = fdf^{-1}$ where the subword f of g is cyclically reduced. Then for any $i \geq 1$ the freely reduced form of g^i is fd^if^{-1}.

Let $n \geq 1$ be the smallest positive integer such that $g^n \in H$. Since $g \notin H$, we have $n > 1$. Suppose that $n > \# V\Gamma$.

Since $g^n \in H$, the freely reduced word fd^nf^{-1} is the label of a reduced path p in Γ from 1_H to 1_H. For each $k = 0, 1, \ldots, n$ let v_k be the end-vertex of the initial segment of p with label fd^k. Since $n > \# V\Gamma$, for some $1 \leq i < j$ we have $v_i = v_j = v$. Then the subpath α of p from v_i to v_j has label d^{j-i}, so that $p = p_1 \alpha p_2$. Put p' be the path obtained from p by removing the subpath α, that is $p' = p_1p_2$. It is obvious that the label of p' is $fd^{n-(j-i)}f^{-1}$. Moreover, since $i \geq 1$, we have $1 \leq n - (j - i) < n$, $n - (j - i) > 0$ and so the word $fd^{n-(j-i)}f^{-1}$ is freely reduced. Since Γ is folded, the path p' is reduced as well.

Therefore by definition of $\Gamma(H)$ we have $g^{n-(j-i)} \in H$ and $1 \leq n - (j - i) < n$. This contradicts the choice of n. □

Corollary 7.11. *There exists an algorithm which, given finitely many freely reduced words*

$$g, h_1, \ldots, h_s \in F(X)$$

decides whether or not some nonzero power of g belongs to the subgroup $H = \langle h_1, \ldots, h_s \rangle$.

The use of subgroup graphs also allows us to easily determine whether one subgroup is conjugate to a subgroup of another.

Lemma 7.12. *Let K, H be subgroups of $F(X)$. Then there is $g \in F(X)$ with $gKg^{-1} \leq H$ if and only if there exists a morphism of (non-based) X-digraphs $\pi : Type(\Gamma(K)) \longrightarrow Type(\Gamma(H))$.*

Proof. Suppose K is conjugate to a subgroup of H. Since conjugation preserves type, we can assume that $\Gamma(H) = Type(\Gamma(H))$. Let $g \in F(X)$ be such that $gKg^{-1} \leq H$. By Proposition 4.3 there exists a morphism $\beta : \Gamma(gKg^{-1}) \longrightarrow \Gamma(H)$. However by Lemma 7.6 $Type(\Gamma(K)) = Type(\Gamma(gKg^{-1}))$ is a subgraph of $\Gamma(gKg^{-1})$. The restriction of β to this subgraph produces the desired morphism $\pi : Type(\Gamma(K)) \longrightarrow Type(\Gamma(H))$.

Suppose now that a morphism $\pi : Type(\Gamma(K)) \longrightarrow Type(\Gamma(H))$ exists. Let f be the label of the shortest path in $\Gamma(K)$ from 1_K to $Type(\Gamma(K))$ and let u be the terminal vertex of this path. Let $v = \pi(u)$ be the image of this vertex in $\Gamma(H)$.

Let c be the label of a shortest path from 1_H to v in $\Gamma(H)$. Thus by Lemma 7.5 $L(Type(K), u) = f^{-1}Kf$ and $L(Type(H), v) = c^{-1}Hc$, so that $(Type(K), u) = (\Gamma(f^{-1}Kf), 1_{f^{-1}Kf})$ and $(Type(H), v) = (\Gamma(c^{-1}Hc), 1_{c^{-1}Hc})$. Since π can also be considered as a morphism $\pi : (\Gamma(f^{-1}Kf), 1_{f^{-1}Kf}) \longrightarrow (\Gamma(c^{-1}Hc), 1_{c^{-1}Hc})$, Proposition 4.3 implies that

$$f^{-1}Kf \leq c^{-1}Hc, \text{ and } K \leq (f \cdot c^{-1})H(f \cdot c^{-1})^{-1},$$

as required. Then K is conjugate to a subgroup of H by Proposition 7.7 and Proposition 4.3. □

Corollary 7.13. *There exists an algorithm which, given finitely many freely reduced words*

$$k_1, \dots, k_s, h_1, \dots, h_m \in F(X),$$

decides whether or not there is $g \in F(X)$ *such that* $gKg^{-1} \le H$, *where* $K = \langle k_1, \dots, k_s \rangle$, $H = \langle h_1, \dots, h_m \rangle$. *Moreover, the algorithm will produce one such* g *if it exists.*

Proof. We first construct the graphs $\Gamma(K), \Gamma(H)$ and their type-subgraphs $Type(\Gamma(K))$, $Type(\Gamma(H))$. Let u be the vertex of $Type(\Gamma(K))$ closest to 1_K (thus $u = 1_K$ if $\Gamma(K) = Type(\Gamma(K))$ and u is the terminal vertex of the segment from 1_K to $Type(\Gamma(K))$ otherwise). Let f be the label of the unique reduced path from 1_K to u in $\Gamma(K)$.

We then enumerate the vertices of $Type(\Gamma(H))$ and for each vertex v check whether there exists a morphism $\pi : Type(\Gamma(K)) \longrightarrow Type(\Gamma(H))$ which takes u to v. If no such vertex of $Type(\Gamma(H))$ is found, K is not conjugate to a subgroup of H by Lemma 7.12. If such v is found, then by the proof of Lemma 7.12 $gKg^{-1} \le H$ where $g = f \cdot c^{-1}$ and c is the label of a reduced path from 1_H to v in $\Gamma(H)$. □

8. Finite Index, Commensurability and Marshall Hall Theorem

Definition 8.1 (Regular digraphs)**.** An X-digraph Γ is said to be X-*regular* if for every vertex v of Γ and every $x \in X \cup X^{-1}$ there is exactly one edge in $\hat{\Gamma}$ with origin v and label x.

We need to prove the following simple graph-theoretic statement.

Lemma 8.2. *Let* Γ *be a finite connected digraph. Let* T *be a spanning tree in* Γ. *Then* $\#E^+(\Gamma - T) = \#E^+\Gamma - \#V\Gamma + 1$.

In particular, if Γ *is a finite connected folded* X*-digraph then* $\#Y_T = \#E^+\Gamma - \#V\Gamma + 1$. *Thus if* v *is a vertex of* Γ *and* $H = L(\Gamma, v)$ *then* $rk(H) = \#E^+\Gamma - \#V\Gamma + 1$.

Proof. It is easy to prove by induction on the number of edges that for a finite tree the number of vertices minus the number of positive edges is equal to one.

Hence $\#E^+T = \#VT - 1 = \#V\Gamma - 1$. Therefore

$$\#E^+\Gamma = \#E^+T + \#E^+(\Gamma - T) = \#V\Gamma - 1 + \#E^+(\Gamma - T)$$

and so

$$\#E^+(\Gamma - T) = \#E^+\Gamma - \#V\Gamma + 1$$

as required. □

The above statement can be used to compute the rank of a finitely generated subgroup $H \le F(X)$ via its graph $\Gamma(H)$.

We can now give a criterion for a subgroup $H \le F(X)$ to be of finite index in $F(X)$.

Proposition 8.3 (Finite index subgroups)**.** *Let* H *be a subgroup of* $F(X)$.

Then $|F(X) : H| < \infty$ *if and only if* $\Gamma(H)$ *is a finite* X*-regular graph. In this case* $|F(X) : H| = \#V\Gamma(H)$.

Proof. Assume first that $|F(X) : H| = s < \infty$. Then H is finitely generated and thus $\Gamma(H)$ is finite. Suppose that $\Gamma(H)$ is not X-regular. Then there is a vertex v of Γ and a letter $x \in X \cup X^{-1}$ such that there is no edge labeled x with origin v in $\hat{\Gamma}$. Choose a reduced path p from 1_H to v in Γ and let $w = \mu(p)$. Then the word wx is freely reduced and, moreover, no word with initial segment wx is accepted by $\Gamma(H)$. Thus for any freely reduced word y whose first letter is not x^{-1}, we have $wxy \notin H$.

Since $|F(X) : H| = s < \infty$, there exist s elements $g_1, \dots, g_s \in F(X)$ such that

$$F(X) = Hg_1 \cup \dots \cup Hg_s$$

Let $M = \max\{|g_i|_X \mid i = 1, \dots, s\}$ and let $f = wx^{M+1}$. Then for some g_i we have $f \cdot g_i^{-1} \in H$. By the choice of M the freely reduced form z of $f \cdot g_i^{-1}$ has initial segment wx and therefore $z \notin L(\Gamma(H), 1_H) = H$. This contradicts our assumption that $f \cdot g_i^{-1} \in H$.

Suppose now that $\Gamma(H)$ is an X-regular finite graph. We want to show that the index of H in $F(X)$ is finite and, moreover, that it is equal to the number of vertices in $\Gamma(H)$.

For each vertex v of $\Gamma(H)$ choose a reduced path p_v from 1_H to v in $\Gamma(H)$ and denote $g_v = \mu(p_v)$. We claim that

$$F(X) = \cup\{H \cdot g_v \mid v \in V\Gamma(H)\}. \tag{1}$$

Let $f \in F(X)$ be an arbitrary freely reduced word. Since $\Gamma(H)$ is X-complete, there is a path α in $\Gamma(H)$ with origin 1_H and label f. Let u be the terminal vertex of this path. Then αp_u^{-1} is a path in $\Gamma(H)$ from 1_H to 1_H with label fg_u^{-1}. Therefore the freely reduced form of the word fg_u^{-1} is accepted by $\Gamma(H)$ and so $f \cdot g_u^{-1} \in H, f \in H \cdot g_u$. Since $f \in F(X)$ was chosen arbitrarily, we have verified that (1) holds.

This already implies that H has finite index in $F(X)$. To prove the proposition it remains to check that if v and u are distinct vertices of $\Gamma(H)$ then $Hg_v \neq Hg_u$. Suppose this is not so and for some $v \neq u$ we have $Hg_v = Hg_u$. Then $g_v \cdot g_u^{-1} = h \in H$. The path $p = p_v p_u^{-1}$ has label $g_v g_u^{-1}$ and hence the path-reduced form p' of p has label h. Since p' is a reduced path starting at 1_H with label $h \in H$, the terminal vertex of p' is also 1_H. Path-reductions do not change the end-points of a path. Therefore the terminal vertex of p is 1_H, as well as the initial vertex of p. However this implies that the paths p_v and p_u have the same terminal vertices and $v = u$, contrary to our assumptions. ☐

Corollary 8.4 (Schrier's Formula). *Let H be a subgroup of finite index i in a finitely generated free group $F = F(X)$. Then*

$$rk(H) - 1 = i(rk(F) - 1)$$

Proof. Let $\Gamma = \Gamma(H)$. Then by Lemma 8.2 $rk(H) - 1 = \#E^+\Gamma - \#V\Gamma$.

Since $|F : H| = i < \infty$, the graph Γ is X-regular and the degree of each vertex is $2\#X = 2rk(F)$. If we add up the degrees of all vertices of Γ, we will count every edge twice. Therefore $2rk(F)\#V\Gamma = 2\#E^+\Gamma$ and $rk(F)\#V\Gamma = \#E^+\Gamma$. Recall also that by Proposition 8.3 $\#V\Gamma = i$. Hence

$$rk(H) - 1 = \#E^+\Gamma - \#V\Gamma = rk(F)\#V\Gamma - \#V\Gamma = rk(F)i - i = i(rk(F) - 1)$$

as required. ☐

Corollary 8.5. *There exists an algorithm which, given finitely many freely reduced words $h_1, \dots, h_s \in F(X)$, computes the index of the subgroup $H = \langle h_1, \dots, h_s\rangle$ in $F(X)$.*

Proof. We first constrict the finite graph $\Gamma(H)$ and then check if it is X-regular. If no, the subgroup H has infinite index in G (by Proposition 8.3). If yes, the index of H in G is equal to the number of vertices of $\Gamma(H)$ (also by Proposition 8.3). ☐

Definition 8.6. Let H be a subgroup of a group G. The *commensurator* $Comm_G(H)$ of H in G is defined as

$$Comm_G(H) = \{g \in G \bigm| |H : H \cap gHg^{-1}| < \infty \text{ and } |gHg^{-1} : H \cap gHg^{-1}| < \infty\}$$

It is easy to see that $Comm_G(H)$ is a subgroup of G containing H.

Lemma 8.7. *Let $H \leq F(X)$ be a nontrivial finitely generated subgroup. Then $|F(X) : H| < \infty$ if and only if $F(X) = Comm_{F(X)}(H)$.*

Proof. It is obvious that $|F(X) : H| < \infty$ implies $F(X) = Comm_{F(X)}(H)$. Suppose now that H is a finitely generated subgroup of $F(X)$ and $F(X) = Comm_{F(X)}(H)$. Assume that $|F(X) : H| = \infty$.

Then the graph $\Gamma(H)$ is not X-regular. Thus there is a vertex $v \in \Gamma$ and a letter $x \in X \cup X^{-1}$ such that there is no edge labeled x with origin v in $\Gamma(H)$. Since H is a nontrivial subgroup of infinite index in $F(X)$, the rank of $F(X)$ is at least two and so $\#(X) \geq 2$. Let $a \in X$ be a letter such that $a \neq x^{\pm 1}$.

Since for any $g \in F(X)$ we have $|gHg^{-1} : H \cap gHg^{-1}| < \infty$, for any element of gHg^{-1} some power of this element belongs to $H \cap gHg^{-1}$ and so to H. Hence for any $g \in F(X)$ and any $h \in H$ there is $n \geq 1$ such that $g^{-1}h^n g \in H$.

Let $h \in H$ be a nontrivial element, so that h is a freely reduced word in X. Let $y \in X \cup X^{-1}$ be the first letter of h and let $z \in X \cup X^{-1}$ be the last letter of h. Then for any $m \geq 1$ the freely reduced form of h^n begins with y and ends with z.

Let w be the label of a reduced path in $\Gamma(H)$ from 1_H to v. Since there is no edge labeled x with origin v in $\Gamma(H)$, any freely reduced word with initial segment wx does not belong to H.

Put $q = y$ if $y \neq z^{-1}$. If $y = z^{-1}$ and $y \in \{x, x^{-1}\}$, put $q = a$. If $y = z^{-1}$ and $y \notin \{x, x^{-1}\}$, put $q = x$. Then for any $m \geq 1$ the word $q\overline{h^m}q^{-1}$ is freely reduced. (Recall that $\overline{h^m}$ is the freely reduced form of h^m.)

Choose a freely reduced word w' such that the word $wxw'q$ is freely reduced. This is obviously possible since X has at least two elements. Put $g = wxw'q$. By our assumptions there is $n \geq 1$ such that the element $g \cdot h^n g^{-1}$ belongs to H. However by the choice of q the freely reduced form of $gh^n g^{-1}$ is $wxw'qy \dots zq^{-1}(w')^{-1}x^{-1}w^{-1}$. The word $wxw'qy \dots zq^{-1}(w')^{-1}x^{-1}w^{-1}$ has initial segment wx and hence cannot represent an element of H. This yields a contradiction. □

Corollary 8.8 (Greenberg-Stallings Theorem)**.** [19], [43] *Let H, K be finitely generated subgroups of $F(X)$ such that $H \cap K$ has finite index in both H and K.*

Then $H \cap K$ has finite index in the subgroup $\langle H \cup K \rangle$.

Proof. Let $G = \langle H \cup K \rangle$ and $L = H \cap K$. Since $|H : L| < \infty$ and $|K : L| < \infty$, we have $H \leq Comm_{F(X)}(L)$, $K \leq Comm_{F(X)}(L)$ and hence $G = \langle H \cup K \rangle \leq Comm_{F(X)}(L)$.

Both H and K are finitely generated and therefore so are G and L (recall that L is a subgroup of finite index in a finitely generated free group H). Hence by Corollary 6.8 G is a free group of finite rank, $G = F(S)$ for some finite set S. Since $G \leq Comm_{F(X)}(L)$, we have $G = Comm_G(L)$. This, by Lemma 8.7, implies that $|G : L| < \infty$, as required. □

Proposition 8.9. *Let H be a nontrivial finitely generated subgroup of $F(X)$. Then H has finite index in $Comm_{F(X)}(H)$.*

Proof. Let $G = Comm_{F(X)}(H)$. If G is finitely generated then G is a free group of finite rank containing H and $G = Comm_G(H)$. Hence $|G : H| < \infty$ by Lemma 8.7.

Suppose now that G is not finitely generated. Then $G = F(Y)$ where Y is an infinite free basis of G. Since H is finitely generated, the elements of a finite generating set for H involve only finitely many letters of Y. Thus H is contained in a finitely generated free factor A of G and $G = A * B$ where B is a free group of infinite rank. However $G = Comm_G(H)$ and so for each $g \in G$ we have $|H : gHg^{-1} \cap H| < \infty$. However for any $b \in B$ $bAb^{-1} \cap A = 1$ and so $bHb^{-1} \cap H = 1$. This contradicts the fact that H is nontrivial and thus infinite. □

Lemma 8.10. *Let Γ be a finite folded X-digraph. Then there exists a finite folded X-regular digraph Γ' such that Γ is a subgraph of Γ' and such that $V\Gamma' = V\Gamma$*

Proof. Let $x \in X$ be a letter. Denote by i_x the number of vertices in Γ which do not have an outgoing edge labelled x. Similarly, denote by j_x the number of vertices in Γ which do not have an incoming edge labelled x. Let n_x be the number of edges in Γ labelled x. It is clear that n_x is equal to the number of vertices having an incoming edge labeled x. Similarly, n_x is equal to the number of vertices having an outgoing edge labeled x. Hence

$$\#V\Gamma - n_x = i_x = j_x.$$

Therefore by adding exactly $i_x = j_x$ edges labelled x to Γ we can obtained a folded X-graph with the same vertex set as Γ where each vertex has an incoming edge labelled x and an outgoing edge labelled x. Repeating this process for each letter $x \in X$ we obtain an X-regular graph Γ' with the properties required in Lemma 8.10. □

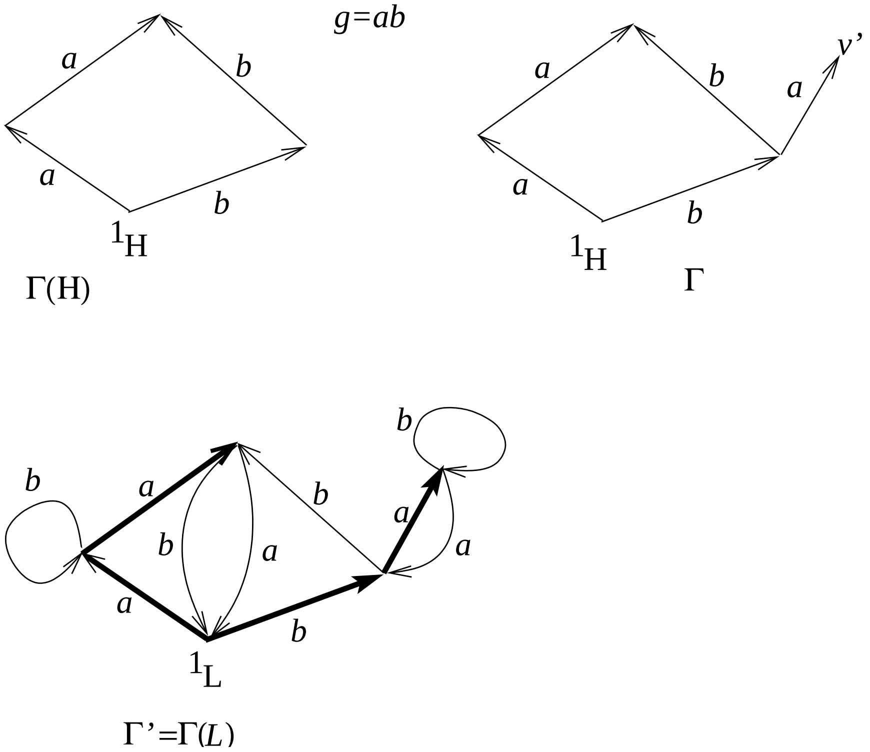

Figure 7. An illustration of M.Hall Theorem. We start with a cyclic subgroup $H = \langle b^2a^{-2}\rangle \le F(a,b)$ (so that $X = \{a,b\}$) and $g = ab \notin H$. Then the graph Γ containing $\Gamma(H)$ is completed to an X-regular graph Γ' corresponding so a subgroup $L \le F(a,b)$ of index 5. The highlighted spanning tree T in $\Gamma' = \Gamma(L)$ can be used to find a free basis of L (containing as a subset a free basis of H): $\{b^2a^{-2}, a^3, a^2b, aba^{-1}, bab^{-1}, ba^2b^{-1}, baba^{-1}b^{-1}\}$. Thus $L = H * K$ where $K = \langle a^3, a^2b, aba^{-1}, bab^{-1}, ba^2b^{-1}, baba^{-1}b^{-1}\rangle$. It is clear that $g = ba \notin L$.

The graph-theoretic criterion of being a subgroup of finite index can also be used to prove the following classical result due to M.Hall [20].

Theorem 8.11 (Marshall Hall's Theorem)**.** *Let H be a finitely generated subgroup of $F(X)$. Let $g \in F(X)$ be such that $g \notin H$.*

Then there exists a finitely generated subgroup K of $F(X)$ such that

1. $L = \langle H, K\rangle = H * K$;
2. *L has finite index in $F(X)$;*
3. $g \notin H$.

Proof. Let $H \le F(X)$ be a finitely generated subgroup. If H has finite index in $F(X)$, then $K = 1$ satisfies the requirement of Theorem 8.11. Suppose now that H has finite index in $F(X)$. Then $\Gamma(H)$ is not X-complete.

Write the element g as a freely reduced word in X and construct an X-graph Z which is a line-segment subdivided into $|g|$ edges with initial vertex u and terminal vertex v such that the label of Z read from u to v is precisely g.

We join the graphs $\Gamma(H)$ and Z by amalgamating 1_H and u into a single vertex. Let Γ_0 be the resulting X-graph. We then perform all possible X-foldings on Γ_0 and denote the resulting graph by Γ. Note that $\Gamma(H)$ is canonically a subgraph of Γ. We can also think about Γ as the result of "wrapping" onto $\Gamma(H)$

the maximal initial segment of g which can be red in $\Gamma(H)$ starting at 1_H. Let v' be the image of v in Γ. Note that $v' \neq 1_H$ since $g \notin H$. The graph Γ is finite, connected and X-folded by construction.

By Lemma 8.10 there exists a connected X-regular graph Γ' containing Γ as a subgraph and having the same vertex set as Γ.

Thus $L = L(\Gamma', 1_H)$ is a subgroup of finite index (equal to $\#V\Gamma$) in $F(X)$. Since $\Gamma(H)$ is a subgraph of Γ and hence of Γ', Proposition 6.2 implies that H is a free factor of L. Moreover, the path labelled g with origin 1_H in Γ' terminates at the vertex $v' \neq 1_H$. Hence $g \notin L$, as required. Theorem 8.11 is proved. ☐

It is easy to see that our proof of M.Hall's Theorem is algorithmically effective. An example of how such an algorithm works is given in Figure 7.

Remark 8.12. Theorem 8.11 implies that a free group $F(X)$ of finite rank is *subgroup separable*, that is to say any finitely generated subgroup H of $F(X)$ is equal to the intersection of finite index subgroups of $F(X)$ containing H. This is an important and nontrivial property of free groups.

Corollary 8.13. *There exists an algorithm which, given a finite set of freely reduced words $h_1, \dots, h_m \in F(X)$ produces the following for the subgroup $H = \langle h_1, \dots, h_m \rangle \leq F(X)$:*

1. *The graph of a finitely generated subgroup $L \leq F(X)$ such that H is a free factor of L and L is a subgroup of finite index in $F(X)$.*
2. *A free basis Y of L decomposed as a disjoint union $Y = Y_H \cup Y_C$, where Y_H is a free basis of H. Thus $L = F(Y) = H * C$ where $C = F(Y_C)$.*
3. *The index of L in $F(X)$.*

Proof. We first construct the graph $\Gamma(H)$ and then complete it to an X-regular folded graph Δ without adding any new vertices. Thus $(\Delta, 1_H) = (\Gamma(L), 1_L)$ and L has finite index in $F(X)$. In fact this index is equal to the number of vertices in $\Gamma(H)$.

We then find a spanning tree T in $\Gamma(H)$. Since the vertex sets of $\Gamma(H), \Gamma(L)$ are the same, T is also a spanning tree in $\Gamma(L)$. We then make a free basis $Y = Y_T$ for L as described in Lemma 6.1. The set Y naturally decomposes into a disjoint union $Y = Y_H \cup Y_C$, where Y_H is a free basis of H (those elements $[e]$ of Y_T corresponding to $e \in E^+(\Gamma(H) - T)$ form Y_H). Thus $L = H * C$ (where $C = F(Y_C)$) as required. ☐

It turns out that X-regular graphs are essential for describing normal subgroups of $F(X)$.

Theorem 8.14 (Normal subgroups)**.** *Let $H \leq F(X)$ be a nontrivial subgroup of $F(X)$. Then H is normal in $F(X)$ if and only if the following conditions are satisfied:*

1. *The graph $\Gamma(H)$ is X-regular (so that there are no degree-one vertices in $\Gamma(H)$ and hence $\Gamma(H) = Type(\Gamma(H))$).*
2. *For any vertex v of $\Gamma(H)$ the based X-digraphs $(\Gamma(H), 1_H)$ and $(\Gamma(H), v)$ are isomorphic (that is $L(\Gamma(H), v) = H$).*

Proof. Suppose first that conditions (1) and (2) are satisfied. Let $g \in F(X)$ be an arbitrary freely reduced word. Since $\Gamma(H)$ is folded and X-regular, there exists a unique path p in $\Gamma(H)$ with origin 1_H and label g. Let v be the terminal vertex of p. By Lemma 7.5 $L(\Gamma(H), v) = gHg^{-1}$. On the other hand by assumption (2) $L(\Gamma(H), v) = H$. Thus $gHg^{-1} = H$. Since $g \in F(X)$ was chosen arbitrarily, this implies that H is normal in $F(X)$.

Suppose now that $H \neq 1$ is normal in $F(X)$. Note that the degree of 1_H in $\Gamma(H)$ is at least two. Indeed, suppose there is only one edge with origin 1_H and let x be the label of this edge. Then any element of H (considered as a freely reduced word) has the form xwx^{-1}. Let $h = xwx^{-1} \in H$ be one such nontrivial element. The cyclically reduced form q of h is conjugate to h and hence belongs to H (since H is normal). However, a cyclically reduced word cannot have the form $x \dots x^{-1}$.

Thus the degree of 1_H is at least two and so $Type(\Gamma(H)) = \Gamma(H)$.

Let $g \in F(X)$ be an arbitrary freely reduced word. We claim that there is a path p_g in $\Gamma(H)$ with origin 1_H and label g. If this is not the case then by Lemma 7.6 the graph of the subgroup gHg^{-1} has

$1_{gHg^{-1}}$ as a degree-one vertex. However, $H = gHg^{-1}$ since H is normal, yielding a contradiction with the above observations about $\Gamma(H)$.

Thus such a path p_g indeed exists. Since $g \in F(X)$ was chosen arbitrarily and the graph $\Gamma(H)$ is folded and connected, this implies that $\Gamma(H)$ is X-regular. Moreover, if v_g is the terminal vertex of p_g then by Lemma 7.5 $L(\Gamma(H), v_g) = gHg^{-1} = H$, so that $(\Gamma(H), v_g) \cong (\Gamma(H), 1_H)$. It is clear that any vertex v of $\Gamma(H)$ can be obtained as v_g for some $g \in F(X)$ and so all the conditions of Theorem 8.14 are satisfied. □

Corollary 8.15. *There exists an algorithm which, given finitely many freely reduced words* $h_1, \dots, h_m \in F(X)$ *decides whether or not* $H = \langle h_1, \dots, h_m \rangle$ *is a normal subgroup of finite index in* $F(X)$.

Proof. We first construct the graph $\Gamma(H)$ and check whether it is X-regular. If not, then the index of H is infinite. If yes, we list the vertices of $\Gamma(H)$ and for each vertex v check if the graphs $(\Gamma(H), 1_H)$ and $(\Gamma(H), v)$ are isomorphic. □

9. Intersections of subgroups

One of the most interesting applications of the folded graphs technique is for computing the intersection of two subgroups of a free group.

Definition 9.1 (Product-graph). Let Γ and Δ be X-digraphs. We define the *product-graph* $\Gamma \times \Delta$ as follows.

The vertex set of $\Gamma \times \Delta$ is the set $V\Gamma \times V\Delta$.

For a pair of vertices $(v, u), (v', u') \in V(\Gamma \times \Delta)$ (so that $v, v' \in V\Gamma$ and $u, u' \in V\Delta$ and a letter $x \in X$ we introduce an edge labeled x with origin (v, u) and terminus (v', u') provided there is an edge, labeled x, from v to v' in Γ and there is an edge, labeled x, from u to u' in Δ.

Thus $\Gamma \times \Delta$ is an X-digraph. In this situation we will sometimes denote a vertex (v, u) of $\Gamma \times \Delta$ by $v \times u$.

The following trivial observation follows immediately from the definition.

Lemma 9.2. *Suppose* Γ *and* Δ *are folded* X*-digraphs. Then* $\Gamma \times \Delta$ *is also a folded* X*-digraph. Moreover, if* $v \times u$ *is a vertex of* $\Gamma \times \Delta$, *then the degree of* $v \times u$ *is less than or equal to the minimum of the degree of* v *in* Γ *and the degree of* u *in* Δ.

Next we establish an important connection between product-graphs and intersections of subgroups.

Lemma 9.3. *Let* Γ *and* Δ *be folded* X*-digraphs. Let* $H = L(\Gamma, v)$ *and* $K = L(\Delta, u)$ *for some vertices* $v \in V\Gamma$ *and* $u \in V\Delta$. *Let* $y = (v, u) \in V(\Gamma \times \Delta)$.

Then $\Gamma \times \Delta$ *is folded and* $L(\Gamma \times \Delta, y) = H \cap K$.

Proof. Let w be a freely reduced word in X.

By definition of $\Gamma \times \Delta$ there exists a path in $\Gamma \times \Delta$ from y to y with label w if and only if there is a reduced path in Γ from v to v with label w and there is a reduced path in Δ from u to u with label w. This implies $L(\Gamma \times \Delta, y) = L(\Gamma, v) \cap L(\Delta, u) = H \cap K$, as required. □

Proposition 9.4. *Let* $H \leq F(X)$ *and* $K \leq F(X)$ *be two subgroups of* $F(X)$. *Let* $G = H \cap K$. *Let* $\Gamma(H) \times_c \Gamma(K)$ *be the connected component of* $\Gamma(H) \times \Gamma(K)$ *containing* $1_H \times 1_K$. *Let* $\Delta = Core(\Gamma(H) \times \Gamma(K), 1_H \times 1_K)$.

Then $(\Gamma(G), 1_G) = (\Delta, 1_H \times 1_K)$.

Proof. By Lemma 9.3 the graph $\Gamma(H) \times_c \Gamma(K)$ is folded and $L(\Gamma(H) \times_c \Gamma(K), 1_H \times 1_K) = H \cap K$. Taking the core does not change the language of a graph, so $L(\Delta, 1_H \times 1_K) = H \cap K = G$.

Obviously, Δ is a folded connected graph which is a core graph with respect to $1_H \times 1_K$. Hence $(\Delta, 1_H \times 1_K) = (\Gamma(G), 1_G)$. □

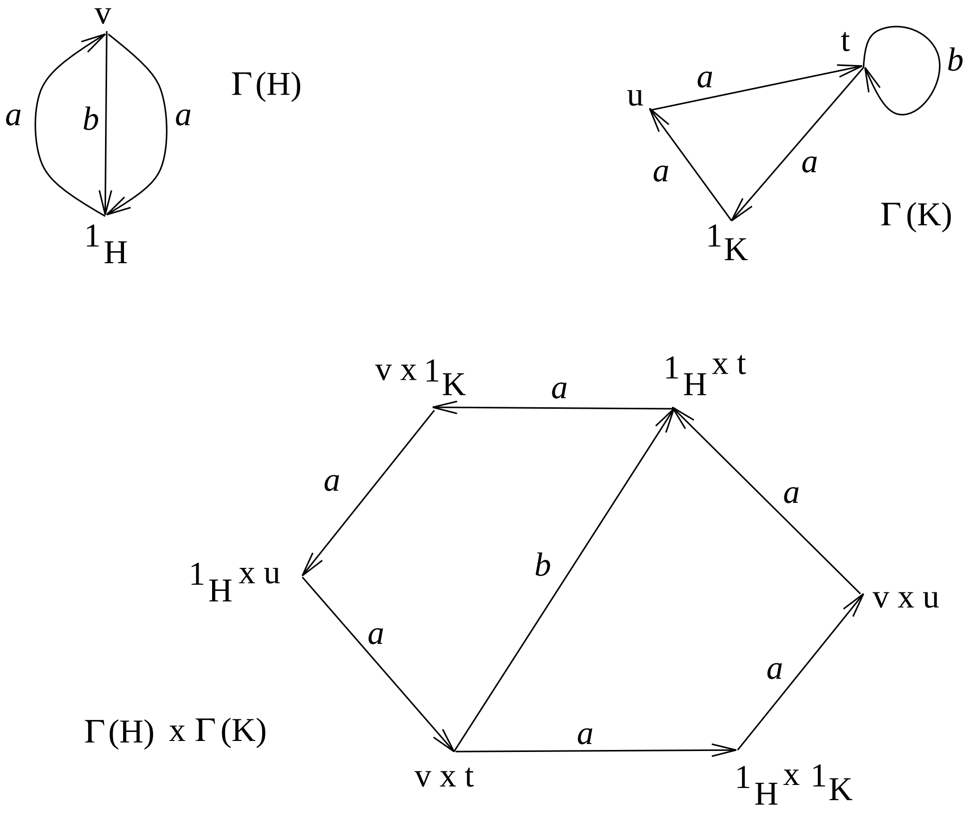

FIGURE 8. Product of two graphs. Here we compute $\Gamma(H) \times \Gamma(K)$ where $H = \langle ab, b^{-1}a\rangle \leq F(a,b)$ and $K = \langle a^3, a^{-1}ba\rangle \leq F(a,b)$. In this example the product is connected and is a core-graph, although in general this need not be the case.

Corollary 9.5. *There exists an algorithm which, given finitely many freely reduced words*

$$h_1, \dots, h_s, k_1, \dots, k_m \in F(X),$$

finds the rank and a Nielsen-reduced free basis of the subgroup $\langle h_1, \dots, h_s\rangle \cap \langle k_1, \dots, k_m\rangle$ *of* $F(X)$. *In particular, this algorithm determines whether or not* $\langle h_1, \dots, h_s\rangle \cap \langle k_1, \dots, k_m\rangle = 1$.

Corollary 9.6 (Howson Property). [22] *The intersection of any two finitely generated subgroups of* $F(X)$ *is again finitely generated.*

Proof. Let H, K be finitely generated subgroups of $F(X)$ and let $G = H \cap K$. By Lemma 5.4 the graphs $\Gamma(H)$ and $\Gamma(K)$ are finite. Therefore the graph $\Gamma(H) \times \Gamma(K)$ is finite. By Lemma 9.3 $G = L(\Gamma(H) \times \Gamma(K), 1_H \times 1_K)$. Therefore by Lemma 6.1 the group G is finitely generated. □

It turns out that even those components of $\Gamma(H) \times \Gamma(K)$ which don't contain $1_H \times 1_K$, carry some interesting information about H and K.

Proposition 9.7. *Let* $H, K \leq F(X)$. *Let* $g \in F(X)$ *be such that the double cosets* KgH *and* KH *are distinct. Suppose that* $gHg^{-1} \cap K \neq 1$.

Then there is a vertex $v \times u$ *in* $\Gamma(H) \times \Gamma(K)$ *which does not belong to the connected component of* $1_H \times 1_K$ *such that the subgroup* $L(\Gamma(H) \times \Gamma(K), v \times u)$ *is conjugate to* $gHg^{-1} \cap K$ *in* $F(X)$.

Proof. Let $g = yz$ where z is the largest terminal segment of g such that z^{-1} is the label of a path in $\Gamma(H)$ with origin 1_H. Denote this path by σ and the terminal vertex of σ by v.

If the word y is not the label of a path in $\Gamma(K)$ with origin 1_K, then by Proposition 7.9 $\langle gHg^{-1}, K\rangle = gHg^{-1} * K$. This implies that $gHg^{-1} \cap K = 1$, contrary to our assumptions.

Thus y is the label of a path τ in $\Gamma(K)$ from 1_K to some vertex u.

By Lemma 7.5 $L(\Gamma(H), v) = zHz^{-1}$ and $L(\Gamma(K), u) = y^{-1}Ky$. Note also that

$$y(zHz^{-1} \cap y^{-1}Ky)y^{-1} = yzHz^{-1}y^{-1} \cap K = gHg^{-1} \cap K$$

and therefore the subgroups $gHg^{-1} \cap K$ and $zHz^{-1} \cap y^{-1}Ky$ are conjugate in $F(X)$.

By Lemma 9.3 we have

$$L(\Gamma(H) \times \Gamma(K), v \times u) = zHz^{-1} \cap y^{-1}Ky \cong gHg^{-1} \cap K,$$

as required.

It remains to show that $v \times u$ does not belong to the connected component of $1_H \times 1_K$ in $\Gamma(H) \times \Gamma(K)$. Suppose this is not the case.

Then there exist a reduced path p_v in $\Gamma(H)$ from 1_H to v and a reduced path p_u in $\Gamma(K)$ from 1_K to u such that their labels are the same, that is $\mu(p_v) = \mu(p_u) = \alpha$.

Note that $p_v\sigma^{-1}$ is a path in $\Gamma(H)$ from 1_H to 1_H and therefore $\overline{\mu(p_v\sigma^{-1})} = \overline{\alpha z} = \alpha \cdot z = h \in H$. Similarly, τp_u^{-1} is a path in $\Gamma(K)$ from 1_K to 1_K and hence $y \cdot \alpha^{-1} = k \in K$. Thus we have

$$g = y \cdot z = y \cdot \alpha^{-1} \cdot \alpha \cdot z = k \cdot h \in KH,$$

contrary to our assumption that $KgH \neq KH$. □

The converse of the above statement also holds:

Proposition 9.8. *Let $H, K \leq F(X)$ be two subgroups of $F(X)$.*

Then for any vertex $v \times u$ of $\Gamma(H) \times \Gamma(K)$ the subgroup $L(\Gamma(H) \times \Gamma(K), v \times u)$ is conjugate to a subgroup of the form $gHg^{-1} \cap K$ for some $g \in F(X)$.

Moreover, if $v \times u$ does not belong to the connected component of $1_H \times 1_K$, then the element g can be chosen so that $KgH \neq KH$.

Proof. Let p_v be a reduced path in $\Gamma(H)$ from 1_H to v with label σ. Similarly, let p_u be a reduced path in $\Gamma(K)$ from 1_K to u with label τ. As we have shown in Lemma 7.5, $L(\Gamma(H), v) = \sigma^{-1}H\sigma$ and $L(\Gamma(K), u) = \tau^{-1}K\tau$. Therefore by Lemma 9.3

$$L(\Gamma(H) \times \Gamma(K), v \times u) = \sigma^{-1}H\sigma \cap \tau^{-1}K\tau \text{ conjugate to } \tau\sigma^{-1}H\sigma\tau^{-1} \cap K,$$

and $g = \tau\sigma^{-1}$ satisfies the requirement of the proposition.

Suppose now that $v \times u$ does not belong to the connected component of $1_H \times 1_K$ in $\Gamma(H) \times \Gamma(K)$ but $g = \tau \cdot \sigma^{-1} \in KH$. Thus $\tau \cdot \sigma^{-1} = kh$ for some $k \in K$, $h \in H$ and therefore

$$k^{-1} \cdot \tau = h \cdot \sigma.$$

Let α be the freely reduced form of the element $k^{-1} \cdot \tau = h \cdot \sigma$. Recall that $k^{-1} \in K$ and so k^{-1} is the label of a reduced p_1 in $\Gamma(K)$ from 1_K to 1_K. Then p_1p_u is a path in $\Gamma(K)$ whose label freely reduces to α. Therefore there is a reduced path p_1' in $\Gamma(K)$ from 1_K to u with label α. Similarly, since $h \in H$, there is a path p_2 in $\Gamma(H)$ from 1_H to 1_H with label h. Hence p_2p_v is a path in $\Gamma(H)$ from 1_H to v whose label freely reduces to α. Again, it follows that there is a reduced path p_2' in $\Gamma(H)$ from 1_H to v with label α. Now the definition of $\Gamma(H) \times \Gamma(K)$ implies that there is a path in $\Gamma(H) \times \Gamma(K)$ from $1_H \times 1_K$ to $v \times u$ with label α. However, this contradicts our assumption that $v \times u$ does not belong to the connected component of $1_H \times 1_K$.

Thus $g \notin KH$ and $KgH \neq KH$, as required. □

The above facts allow us to use the product-graphs to decide whether a finitely generated subgroup of a free group is *malnormal*. .

Definition 9.9. Let H be a subgroup of a group G. We say that H is a *malnormal* subgroup of G if for any $g \in G - H$

$$gHg^{-1} \cap H = 1.$$

We say that H is *cyclonormal* if for any $g \in G - H$

$$gHg^{-1} \cap H \text{ is cyclic.}$$

(The notion of a cyclonormal subgroup was suggested by D.Wise [49].)

Malnormal subgroups of free groups and word-hyperbolic groups have recently become the object of intensive studies. In particular, malnormality plays an important role in the Combination Theorem for hyperbolic groups [4], [5], [17], [29]. Various examples, where malnormal subgroups play an important role, can be found in [28], [2], [23], [24], [26], [27], [47], [48], [49], [50], [8] and other sources.

Theorem 9.10. *Let $H \le F(X)$ be a subgroup. Then H is malnormal in $F(X)$ if and only if every component of $\Gamma(H) \times \Gamma(H)$, which does not contain $1_H \times 1_H$, is a tree.*

Proof. Suppose H is malnormal in G. Assume that there is a connected component C of $\Gamma(H) \times \Gamma(H)$, which does not contain $1_H \times 1_H$ and which is not a tree. Let $v \times u$ be a vertex of C. Since C is not a tree, there exists a nontrivial reduced path from $v \times u$ to $v \times u$ in C and hence $A = L(C, v \times u) \neq 1$. However by Proposition 9.8 there is $g \in F(X)$ such that $HgH \neq HH = H$ (so that $g \notin H$) and such that $gHg^{-1} \cap H$ is conjugate to A and thus $gHg^{-1} \cap H \neq 1$. This contradicts our assumption that H is malnormal.

Since for any tree T and any vertex v of T $L(T, v) = 1$, the opposite implication of the theorem is just as obvious and follows immediately from Proposition 9.7. □

Corollary 9.11. *There exists an algorithm which, given finitely many freely reduced words $h_1, \dots, h_s$ in $F(X)$ decides whether or not the subgroup $H = \langle h_1, \dots, h_s \rangle$ is malnormal in $F(X)$. If H is not malnormal, this algorithm will produce a nontrivial element $g \in F(X) - H$ such that $gHg^{-1} \cap H \neq 1$.*

Proof. We first construct the finite graph $\Gamma(H)$ as described in Proposition 7.1 and then build the finite graph $\Gamma(H) \times \Gamma(H)$. It remains to check whether those connected components of $\Gamma(H) \times \Gamma(H)$ which do not contain $1_H \times 1_H$, are trees.

If all these components are trees, H is malnormal in G. If some component Ω not containing $1_H \times 1_H$ is not a tree, then H is not malnormal in G. Moreover, in this case we choose a vertex $(v, u) \in \Omega$ and a paths with labels σ, τ in $\Gamma(H)$ from 1_H to v, u respectively. Put $g = \tau\sigma^{-1}$. Then by the proof of Proposition 9.8 we have $gHg^{-1} \cap H \neq 1$ and $HgH \neq H1H$, so that $g \notin H$. □

A different proof of the above result was given in [2].

Definition 9.12 (Immersed subgroup)**.** Let $H = \langle h_1, \dots, h_s \rangle \le F(X)$ be a finitely generated subgroup of $F(X)$, where each h_i is a nontrivial freely reduced word.

We say that H is *immersed* in G if for any $1 \le i, j \le s$ we have

$$|h_i \cdot h_j|_X = |h_i|_X + |h_j|_X$$

and for any $i \neq j, 1 \le i, j \le s$ we have

$$|h_i \cdot h_j^{-1}|_X = |h_i|_X + |h_j|_X$$

Remark 9.13. Note that if H is an immersed subgroup then for any freely reduced word $w(y_1, \dots, y_s)$ in the alphabet $\{y_1^{\pm 1}, \dots, y_s^{\pm 1}\}$ the word W, obtained from w by substituting y_i by h_i, is freely reduced in $X \cup X^{-1}$. It is also easy to see that for $i \neq j$ the words h_i and h_j have different first letters and different last letters. Moreover, each h_i is cyclically reduced, that is the first letter of h_i is not the inverse of the last letter of h_i. For this reason the map $\sigma : \{\pm 1, \dots, \pm s\} \longrightarrow X \cup X^{-1}$ defined as

$$\sigma : i \mapsto \text{ the first letter of } h_{|i|}^{sign(i)}$$

is injective. Therefore $s \le \#(X)$.

There are many examples of immersed subgroups (see [25] for details). Recall that $X = \{x_1, \ldots, x_N\}$. For each $i, 1 \le i \le n$ let h_i be a nontrivial freely reduced word in X such that the first and the last letters of h_i are x_i. Then the subgroup $H = \langle h_1, \ldots, h_n\rangle$ is obviously immersed in $F(X)$.

The following theorem was first proved, in the context of Stallings' original approach, by D.Wise [50]. An interesting application of this result was given by I.Kapovich [25].

Theorem 9.14. *Let $H = \langle h_1, \ldots, h_s\rangle \le F(X)$ be an immersed subgroup of $F(X)$. Then H is cyclonormal in $F(X)$.*

Proof. Consider the graph Γ which is a wedge of s circles with labels $h_1, \ldots, h_s$ joined at a single vertex (which we denote 1_H). The fact that H is immersed in $F(X)$ implies (and, moreover, is equivalent to the condition) that Γ is already a folded X-digraph. Hence $\Gamma = \Gamma(H)$.

We know from Proposition 9.7 that for any $g \notin H$ the subgroup $gHg^{-1} \cap H$ is either trivial or isomorphic to $L(C, v \times u)$ for some connected component C of $\Gamma(H) \times \Gamma(H)$ which does not contain $1_H \times 1_H$.

Thus to prove cyclonormality of H it is enough to establish that for each such component C and any vertex $v \times u$ of C the subgroup $L(C, v \times u)$ is cyclic (note trivial group is cyclic).

Note that by Lemma 9.2 the degree of a vertex $a \times b$ in a product-graph $A \times B$ is bounded by the minimum of the degrees of a and b in A and B respectively. As we observed before, the degree of every vertex, other than 1_H in $\Gamma(H)$ is equal to two. Hence for each vertex $v \times u \ne 1_H \times 1_H$ in $\Gamma(H) \times \Gamma(H)$ the degree of $v \times u$ is at most two. Thus C is a finite connected graph where the degree of each vertex is at most two. Therefore C is either a segment or a circle. In the first case the language of C (with respect to any of its vertices) is trivial and in the second case this language is an infinite cyclic subgroup of $F(X)$. This implies that H is cyclonormal in $F(X)$. ☐

Corollary 9.15. *There is an algorithm which, given finitely many nontrivial freely reduced words*

$$h_1, \ldots, h_s \in F(X),$$

decides whether or not the subgroup $H = \langle h_1, \ldots, h_s\rangle \le F(X)$ is cyclonormal.

Proof. If follows from Proposition 9.7 that if $g \notin H$ then $gHg^{-1} \cap H$ is either trivial or isomorphic to $L(C, v \times u)$ where C is a connected component of $\Gamma(H) \times \Gamma(H)$ not containing $1_H \times 1_H$ and where $v \times u$ is a vertex of C. Recall also that the isomorphism type of $L(C, v \times u)$ depends only on C and not on the choice of $v \times u$ in C.

Thus to check whether H is cyclonormal we construct $\Gamma(H)$, $\Gamma(H) \times \Gamma(H)$ and for every connected component C of $\Gamma(H) \times \Gamma(H)$, not containing $1_H \times 1_H$ we find a spanning tree T_C in C and count the number n_C of positive edges in $C - T_C$ (this number is equal to the rank of the free group $L(C, v \times u)$). If for each C $n_C \le 1$, the subgroup H is cyclonormal. If for some C we have $n_C \ge 2$, the subgroup H is not cyclonormal. ☐

10. Principal quotients and subgroups

In this section we want to investigate more carefully the relationship between $\Gamma(H)$ and $\Gamma(K)$ if $K \le H$. To simplify the matters we will assume that both H and K are finitely generated, although it will be clear that one can consider a more general situation.

Definition 10.1. Let $K \le H \le F(X)$ be two finitely generated subgroups of $F(X)$. Since $K = L(\Gamma(K), 1_K) \subseteq L(\Gamma(H), 1_H) = H$, for any path p in $\Gamma(K)$ from 1_K to 1_K there is a unique path $\alpha[p]$ in $\Gamma(H)$ from 1_H to 1_H such that the label of $\alpha[p]$ is the same as that of p.

Define a subgraph $\Gamma_H(K)$ of $\Gamma(H)$ as

$$\Gamma_H(K) = \cup\{\alpha[p] \,|\, p \text{ is a path in } \Gamma(K) \text{ from } 1_K \text{ to } 1_K\}.$$

In fact, constructing $\Gamma_H(K)$ can be considerably simplified, as the following lemma shows.

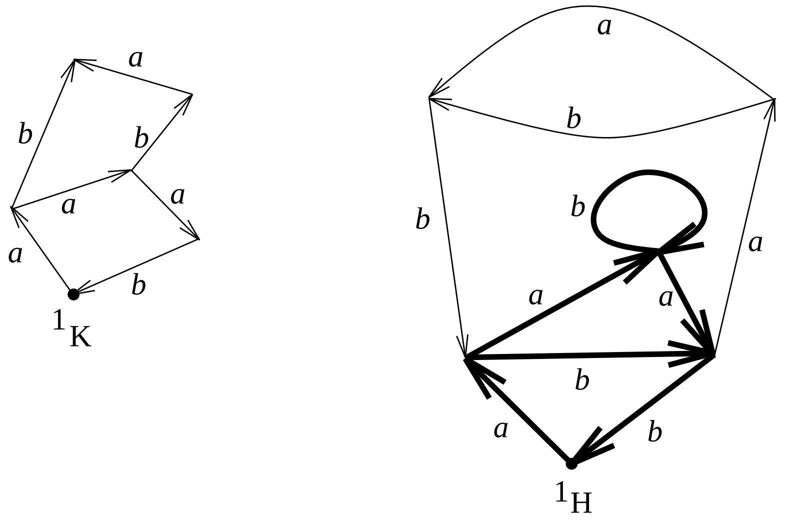

FIGURE 9. An example of $\Gamma_H(K)$. Here $K \leq H \leq F(a,b)$ and the graph $\Gamma_H(K)$ is highlighted. We suggest that the reader uses Lemma 10.2 to verify the validity of this picture

Lemma 10.2. *Let $K \leq H \leq F(X)$ be finitely generated subgroups of $F(X)$. Let $K = \langle k_1, \dots, k_m \rangle$ where k_i are freely reduced words in $F(X)$. For each $i = 1, \dots, m$ let p_i be the loop at 1_K in $\Gamma(K)$ with label k_i. Then*

$$\Gamma_H(K) = \cup_{i=1}^{n} \alpha[p_i]$$

The above lemma shows that in practical terms it is very easy to construct $\Gamma_H(K)$, given $\Gamma(K)$ and $\Gamma(H)$. For example, we can choose as a finite generating set for K the set T_K corresponding to some spanning tree in $\Gamma(K)$.

Lemma 10.3. *Let $K \leq H \leq F(X)$. Then there is a unique morphism of based X-digraphs $\alpha : (\Gamma(K), 1_K) \longrightarrow (\Gamma(H), 1_H)$; for this morphism we have $\alpha(\Gamma(K)) = \Gamma_H(K)$.*

Proof. The existence and uniqueness of α follow from Proposition 4.3.

It is obvious from the definition of $\Gamma_H(K)$ and the fact that any edge in $\Gamma(K)$ lies on a reduced path from 1_K to 1_K that $\alpha(\Gamma(K)) = \Gamma_H(K)$. ☐

Lemma 10.4. *Let $K \leq H \leq F(X)$. Then the graph $\Gamma_H(K)$ is a connected folded graph which is a core graph with respect to 1_H.*

Proof. It is obvious that $\Gamma_H(K)$ is folded, connected and that it contains 1_H. Let $\alpha : \Gamma(K) \longrightarrow \Gamma(H)$ be the unique morphism such that $\alpha(1_K) = 1_H$, so that $\alpha(\Gamma(K)) = \Gamma_H(K)$ (the existence of α follows from Lemma 10.3.

Assume that $\Gamma_H(K)$ is not a core-graph with respect to 1_H and that there is a vertex $u \neq 1_H$ of degree one on $\Gamma_H(K)$. Let e' be the unique edge of $\Gamma_H(K)$ with terminus u. Let e be an edge of $\Gamma(K)$ such that $\alpha(e) = e'$. Since $\Gamma(K)$ is a core graph with respect to 1_K, there exists a reduced loop p at 1_K passing through p. Denote the label of p by w. Since p is reduced and $\Gamma(K)$ is folded, w is freely reduced. The path $\alpha(p)$ is a loop at 1_H in $\Gamma_H(K)$ with label w and $\alpha(p)$ passes through e'. Since w is freely reduced and $\Gamma_H(K)$ is folded, the path $\alpha(p)$ is reduced. However, this contradicts our assumption that u has degree one in $\Gamma_H(K)$. ☐

Recall that by Lemma 4.4 (morphism factorization) an epimorphic image of an X digraph Γ can be obtained by first identifying some vertex subsets of Γ into single vertices and then folding some edges with the same initial vertices and the same terminal vertices (so that the vertex set of the graph does not change at this step).

Definition 10.5 (Principal quotients)**.** Let K be a finitely generated subgroup of $F(X)$. Since $\Gamma(K)$ is finite, there exist only finitely many based X-digraphs $(\Gamma_1, v_1), \dots (\Gamma_s, v_s)$ such that:
(a) each Γ_i is a finite connected folded X-digraph which is a core-graph with respect to v_i;
(b) for each $i = 1, \dots, s$ there is an epimorphism $f : \Gamma(K) \longrightarrow \Gamma_i$ with $f(1_K) = v_i$.

These graphs $(\Gamma_1, v_1), \dots (\Gamma_s, v_s)$ are called *principal quotients* of $\Gamma(K)$.

Remark 10.6. Note that if (Γ, v) is a principal quotient of $\Gamma(K)$ then $(\Gamma, v) = (\Gamma(H), 1_H)$ for some finitely generated subgroup H of $F(X)$ with $K \le H$. Moreover, in this case $\Gamma_H(K) = \Gamma$.

Lemma 10.7. *Let F be a free group (of possibly infinite rank) and let K be a finitely generated free factor of F. Let $H \le F$ be a finitely generated subgroup such that $K \le H$. Then K is a free factor of H.*

Proof. Let $Y = Y_1 \sqcup Y_2$ be a free basis of F such that Y_1 is a finite free basis of K. Since H is finitely generated and involves only finitely many letters from Y_2, we may assume that Y_2 is also finite. It is clear that in this case the graph $\Gamma^Y(K)$ is just a wedge of circles, labeled by elements of Y_1, wedged at a vertex 1_K. This graph clearly has no epimorphic images other than itself, so $\Gamma^Y(K) = \Gamma^Y_H(K) \subseteq \Gamma^Y(H)$. Since $(\Gamma^Y(K), 1_K)$ is a subgraph of $(\Gamma^Y(H), 1_H)$, this implies that K is a free factor of H, as required. □

11. Algebraic extensions

In this section we will develop an elementary analogue of the field extension theory for subgroups of free groups.

Definition 11.1 (Free and algebraic extensions)**.** Let $K \le H \le F(X)$. In this case we will say that H is an *extension* of K. We say that H is a *free extension* of K if there is a subgroup $C \le F(X), C \ne 1$ and a subgroup K' of $F(X)$ such that $K \le K'$ and $H = K * C$.

If an extension $K \le H$ is not free, we call it *algebraic*.

This definition is motivated by the analogy with field extensions. Indeed, if $K \le K'$ and $H = K' * C$ then H is obtained by adding to K' several "purely transcendental" elements, namely a free basis of C. That is why it is reasonable to call such extensions $K \le H$ "free".

Theorem 11.2.
1. *Let $K \le F(X)$ be a finitely generated subgroup. If $K \le H$ (where $H \le F(X)$) is an algebraic extension then $(\Gamma(H), 1_H)$ is a principal quotient of $\Gamma(K)$. In particular, H is finitely generated.*
2. *For a finitely generated $K \le F(X)$ there are only finitely many subgroups $H \le F(X)$ such that $K \le H$ is an algebraic extension.*
3. *Let $K \le H$ be finitely generated subgroups of $F(X)$ and suppose that the extension $K \le H$ is free. Then there exist subgroups K', C such that $K \le K'$, $H = K' * C$ and $(\Gamma(K'), 1_{K'})$ is a principal quotient of $\Gamma(K)$.*
4. *Let $K \le H$ be finitely generated subgroups of $F(X)$. Then there is a free factor K' of H such that $K \le K'$ is an algebraic extension.*

Proof. (1) Let $K \le H$ (where $H \le F(X)$) be a algebraic extension. Then there is a canonical morphism $\alpha : \Gamma(K) \longrightarrow \Gamma(H)$ such that $\alpha(1_K) = 1_H$ and $\alpha(\Gamma(K)) = \Gamma_H(K)$. Since K is finitely generated, the graphs $\Gamma(K)$ and $\Gamma_H(K)$ are finite. Denote $K' = L(\Gamma_H(K), 1_H)$. Then K' is a finitely generated subgroup containing K. Moreover, since $\Gamma_H(K)$ is a subgraph of $\Gamma(H)$, the subgroup K' is a free factor of H. That is, $H = K' * C$. Since the extension $K \le H$ is algebraic, $C = 1$ and so $H = K'$. Since $\Gamma_H(K)$ is a connected subgraph of $\Gamma(H)$ which is a core graph with respect to 1_H and such that $L(\Gamma_H(K), 1_H) = H$, we have $\Gamma_H(K) = \Gamma(H)$. Thus $(\Gamma(H), 1_H)$ is a principal quotient of $\Gamma(K)$.
(2) If K is a finitely generated subgroup of $F(X)$ then $\Gamma(K)$ is finite and hence has only finitely many principal quotients. Therefore by (1) there are only finitely many algebraic extensions of K in $F(X)$.

(3) Let $K \le H$ be finitely generated subgroups of $F(X)$. Let K', C be such that $H = K' * C$, $K \le K'$ and C is of maximal possible rank.

We claim that $\Gamma(K', 1_{K'})$ is a principal quotient of $\Gamma(K)$. Indeed, since $K \le K'$, we can find the subgraph $\Gamma_{K'}(K)$ of $\Gamma(K')$, which is a principal quotient of $\Gamma(K)$. If $\Gamma_{K'}(K)$ is a proper subgraph in $\Gamma(K')$ then by Proposition 6.2 the subgroup $D = L(\Gamma_{K'}(K), 1_{K'})$ is a proper free factor of K' and $K' = D * D'$ with $D' \ne 1$. On the other hand $K \le D$ by construction. Thus $K \le D$ and $H = K' * C = (D * D') * C = D * (D' * C)$. Since the rank of $D' * C$ is greater than the rank of C, we obtain a contradiction with the choice of C.

Thus $\Gamma_{K'}(K) = \Gamma(K')$, and therefore $(\Gamma(K'), 1'_K)$ is indeed a principal quotient of $\Gamma(K)$.

(4) Let $K \le H$ be finitely generated subgroups of $F(X)$.

Choose a free factor K' of H of the smallest possible rank such that K' contains K. Thus $H = K' * C$ for some (possibly trivial) $C \le H$.

We claim that $K \le K'$ is an algebraic extension. Indeed, suppose not. Then there exist nontrivial groups K'', C' such that $K \le K''$ and $K' = K'' * C'$. This means that $H = K'' * (C' * C)$ and $K \le K''$. Moreover $C' \ne 1$ and $K' = K'' * C'$ imply that $rk(K'') < rk(K')$. Thus K'' is a free factor of H containing K and the rank of K'' is smaller than the rank of K'. This contradicts the choice of K'. ☐

Theorem 11.3 (Decidability of extension type)**.** 1. *There is an algorithm which, given finitely many freely reduced words $k_1, \dots, k_m$ in $F(X)$ finds all principal quotients of $\Gamma(K)$, where $K = \langle k_1, \dots, k_m \rangle$.*

2. *There is an algorithm which, given finitely many freely reduced words $k_1, \dots, k_m, h_1, \dots, h_s$ in $F(X)$ decides whether or not the subgroup $H = \langle h_1, \dots, h_s \rangle$ is an algebraic extension of the subgroup $K = \langle k_1, \dots, k_m \rangle$.*
3. *There is an algorithm which, given finitely many elements $f_1, \dots, f_s \in F(X)$, finds all possible algebraic extensions of the subgroup $\langle f_1, \dots, f_s \rangle$ in $F(X)$.*

Proof. (1) First we construct the graph $\Gamma(K)$, as described in Proposition 7.1. We compute all the principal quotients $(\Gamma_1, v_1), \dots, (\Gamma_r, v_r)$ of $\Gamma(K)$ as follows. We list all possible partitions of the vertex set of $\Gamma(K)$ into finitely many disjoint subsets. For each partition we identify each of this subsets into a single vertex. We then perform all possible foldings on the resulting graph which do not change the number of vertices. If the final graph is a folded graph which is a core graph with respect to the image of 1_K, it is a principal quotient of $\Gamma(K)$. Lemma 4.4 ensures that all principal quotients of $\Gamma(K)$ can be obtained in such a way.

(2) First we construct $\Gamma(H)$ and $\Gamma(K)$ and check whether K is contained in H. If not, then H is not an extension of K and thus not an algebraic extension of K.

Suppose $K \le H$. We then find all the principal quotients $(\Gamma_1, v_1), \dots, (\Gamma_r, v_r)$ of $\Gamma(K)$, as described in part (1). After that, for each $i = 1, \dots, r$ we apply the algorithm of Whitehead (see [32], [33]) to check whether $K_i = L(\Gamma_i, v_i)$ is a proper free factor of H. If for some i the answer is yes, then obviously $K \le K_i$, K_i is a proper free factor of H and thus $K \le H$ is a free extension.

If for each $i = 1, \dots, r$ the subgroup K_i is not a proper free factor of H, then by part (3) of Theorem 11.2 the extension $K \le H$ is algebraic.

(3) We first construct the graph $\Gamma(K)$ and all of its principal quotients $(\Gamma_1, v_1), \dots, (\Gamma_r, v_r)$. Denote $H_i = L(\Gamma_i, v_i)$ for $i = 1, \dots, r$, so that $K \le H_i$ for each i. By part (1) of Theorem 11.2 if $K \le H$ is an algebraic extension then $(\Gamma(H), 1_H)$ is a principal quotient of $\Gamma(K)$.

Thus it suffices to check for each $i = 1, \dots, r$ whether $K \le H_i$ is algebraic. The set of those H_i, for which the answer is "yes", gives us all algebraic extensions of K. ☐

Remark 11.4. The proof of Theorem 11.3 is one of the few places in this paper where the algorithm we provide is slow. Primarily this is because we need to use the Whitehead algorithm for deciding whether one subgroup is a free factor of another.

Lemma 11.5. *Let K be a finitely generated subgroup of $F(X)$. Suppose $X = X_1 \cup X_2$ and suppose that $K \le F(X_1)$ (that is the graph $\Gamma(K)$ does not have any edges labeled by elements of $X_2^{\pm 1}$). Let $H \le F(X)$ be a subgroup containing K such that the extension $K \le H$ is algebraic.*

Then $H \le F(X_1)$ *(so that the graph* $\Gamma(H)$ *does not have any edges labeled by elements of* $X_2^{\pm 1}$*).*

Theorem 11.6. *Let* K *be a finitely generated subgroup of* $F(X)$*.*

1. *If* $K \le H \le F(X)$ *and* $|H : K| < \infty$ *then* $K \le H$ *is an algebraic extension.*
2. *If* $K \le H$ *and* $H \le Q$ *are algebraic then* $K \le Q$ *is algebraic.*
3. *If* $K \le H_1$ *and* $K \le H_2$ *are algebraic then* $K \le H$ *is also algebraic, where* $H = \langle H_1, H_2 \rangle$*.*

Proof. (1) This statement follows directly from Definition 11.1.
(2) Suppose the statement fails and $K \le Q$ is a free extension. Note that since K is finitely generated, $K \le H$ is algebraic and $H \le Q$ is algebraic, both H and Q are also finitely generated. Let$Q = K' * C$ where $K \le K'$ and $C \ne 1$. Let Y' be a free basis of K' and let Y be a free basis of C. Then $Z = Y' \cup Y$ is a free basis of Q, so that $Q = F(Z)$. Since $K \le K' = F(Y')$, the graph $\Gamma^Z(K)$ (corresponding to K with respect to the basis Z of Q) does not have any edges labeled by letters of $Y^{\pm 1}$. Since $K \le H$ is algebraic, Lemma 11.5 implies that the graph $\Gamma^Z(H)$ also does not have any edges labeled by elements of $Y^{\pm 1}$. This means that $H \le F(Y')$. Since $Q = F(Y') * F(Y)$ and $Y \ne \emptyset$, the extension $H \le Q$ is free. However, this contradicts our assumption that $H \le Q$ is algebraic.
(3) Suppose the statement fails and $K \le H$ is a free extension. Note that K, H_1, H_2 and H are finitely generated. Thus $H = K' * C$ where $K' = F(Y')$ and $C = F(Y)$ with finite Y, Y' and nonempty Y. Denote $Z = Y' \cup Y$, so that $H = F(Z)$. Since $K \le K' = F(Y')$, the graph $\Gamma^Z(K)$ does not have any edges with labels from $Y^{\pm 1}$. Since $K \le H_1$ and $K \le H_2$ are algebraic, Lemma 11.5 implies that $H_1 \le F(Y')$ and $H_2 \le F(Y')$. Therefore $H = \langle H_1, H_2 \rangle \le F(Y')$. However, this contradicts our assumption that $H = F(Y') * F(Y)$ with $Y \ne \emptyset$. □

Definition 11.7 (Algebraic closure)**.** Let $K \le F(X)$ be a finitely generated subgroup of $F(X)$. Since there are only finitely many algebraic extensions of K in $F(X)$, Theorem 11.6 implies that there exists the largest algebraic extension H of K (which contains all other such extensions). We call this H *the algebraic closure of* K and denote $H = cl(K)$ (or $H = cl_{F(X)}(K)$). If $K = cl(K)$ we say that K is an *algebraically closed* subgroup of $F(X)$.

Lemma 11.8. *There is an algorithm which, given finitely many freely reduced words* $h_1, \dots, h_m \in F(X)$ *finds a free basis of the subgroup* $cl(H)$ *of* $F(X)$ *where* $H = \langle h_1, \dots, h_m \rangle$*.*

Proof. This follows directly from decidability of the extension type and the fact that we can effectively find all algebraic extensions of H in $F(X)$ (Theorem 11.3). □

It turns out that algebraically closed subgroups of $F(X)$ are precisely the free factors of $F(X)$.

Lemma 11.9. *Let* $K \le H$ *be finitely generated subgroups of* $F(X)$ *such that* H *is a free extension of* K*. Let* $C \ne 1$ *and* K' *be such that* $K \le K'$*,* $H = K' * C$ *and* C *is of the largest possible rank. Then* $K \le K'$ *is an algebraic extension.*

Proof. Suppose the statement of Lemma 11.9 fails and the extension $K \le K'$ is free. This means that $K \le K''$ and $K' = K'' * C''$ where $C'' \ne 1$. Therefore

$$H = K' * C = (K'' * C'') * C = K'' * (C'' * C).$$

Since $C'' \ne 1$, $rk(C'' * C) = rk(C'') + rk(C) > rk(C)$, which contradicts the choice of C. □

Theorem 11.10 (Algebraically closed subgroups)**.** *Let* K *be a finitely generated subgroup of* $F(X)$*. Then* K *is algebraically closed in* $F(X)$ *if and only if* H *is a free factor of* $F(X)$*, that is* $F(X) = H * C$ *for some* $C \le F(X)$*.*

Proof. First assume that K is a free factor of $F(X)$. Without loss of generality we may assume that $K = F(X_1)$ where $X = X_1 \cup X_2$. Suppose there $K \le H \le F(X)$ where $K \le H$ is an algebraic extension. Hence by Lemma 11.5 $H \le F(X_1) = K$, so that $H = K$. Hence K is in fact algebraically closed, as required.

Suppose now that K is algebraically closed and $K \ne F(X)$. Therefore the extension $K \le F(X)$ is free. Let K', C be such that $K \le K'$, $F(X) = K' * C$ and $C \ne 1$ is of the biggest rank possible. We claim that $K = K'$. Indeed, if $K \neq K'$ then the extension $K \leq K'$ is free since K is algebraically closed. Thus there is $C' \neq 1$ and K'' such that $K \leq K''$ and $K' = K'' * C'$. This means that

$$F(X) = K' * C = (K'' * C') * C = K'' * (C' * C)$$

where $K \leq K''$ and the rank $rk(C' * C) = rk(C') + rk(C)$ is larger than the rank of C. This contradicts the choice of C. Thus $K = K'$ so that $F(X) = K * C$ and K is a free factor of $F(X)$, as required. ☐

12. A REMARK ON HANNA NEUMANN'S CONJECTURE

By a well-known result of A.Howson (see Corollary 9.6 above) the intersection of any two finitely generated subgroups in a free group is again finitely generated. However, it is interesting to investigate in more detail the connection between the rank of the intersection and the ranks of the two intersecting subgroups. A very important and still open classical conjecture regarding this question was proposed by Hanna Neumann [34]:

Conjecture 12.1. *Let A, B be finitely generated subgroups of a free group $F(X)$. Let $C = A \cap B$ and suppose that $C \neq 1$. Then*

$$rk(C) - 1 \leq (rk(A) - 1)(rk(B) - 1).$$

Over the last four decades the Hanna Neumann conjecture has been the subject of extensive research and a great deal is known by now (see for example [12], [35], [10]), although the conjecture itself remains open. The following simple statement says that in order to prove the conjecture it is enough to establish its validity for the case when both A and B are algebraic extensions of $C = A \cap B$.

Theorem 12.2. *Suppose the Hanna Neumann conjecture holds for all finitely generated subgroups A and B of a free group $F(X)$ such that $C = A \cap B \neq 1$ and such that both A and B are algebraic extensions of $C = A \cap B$. Then the Hanna Neumann conjecture holds for all finitely generated subgroups A and B of $F(X)$ with nontrivial intersection.*

Proof. Suppose the Hanna Neumann conjecture holds for all finitely generated subgroups of $F(X)$ such that both of them are algebraic extensions of their intersection.

Let A and B be two arbitrary finitely generated subgroups of $F(X)$. Let $C = A \cap B$ and suppose that $C \neq 1$. By part (4) of Theorem 11.2 there exist subgroups A', B' such that:
(1) we have $C \leq A' \leq A$, $C \leq B' \leq B$;
(2) the subgroup A' is a free factor of A and the subgroup B' is a free factor of B;
(3) both A' and B' are algebraic extensions of C.

Note that obviously $C = A' \cap B'$. Observe also that (2) implies that $rk(A') \leq rk(A)$ and $rk(B') \leq rk(B)$. By our assumption the Hanna Neumann conjecture holds for algebraic extensions. Therefore

$$rk(C) - 1 \leq (rk(A') - 1)(rk(B') - 1) \leq (rk(A) - 1)(rk(B) - 1).$$

Since A and B were chosen arbitrarily, this implies the statement of Theorem 12.2. ☐

13. MALNORMAL CLOSURES AND ISOLATORS

We have already mentioned the importance of malnormal subgroups. Recently the significance of *isolated* subgroups has also been clarified. Recall that a subgroup H of a group G is called *isolated* in G if whenever $g^n \in H, g \in G, n > 0$, we also have $g \in H$. It turns out that isolated subgroups (in particular those of free groups) play a substantial role in the study of equations over free groups and the elementary theory of free groups (see [31], [30]). In this section we obtain some new results regarding malnormal and isolated subgroups of free groups.

Unlike malnormality (see Theorem 9.10), verifying the property of being isolated is not as simple. Nonetheless, it can still be done using the subgroup graph.

Theorem 13.1. *There exists an algorithm which, given a finite collection of words $h_1, \dots, h_l \in F(X)$ decides if the subgroup $H = \langle h_1, \dots, h_l \rangle$ is isolated in $F(X)$.*

Proof. Let $\Gamma = \Gamma(H)$ be the subgroup graph of H. Denote $n := \#X$ and $k := \#V\Gamma$.

We will first prove the following

Claim. Suppose $g^m \in H$ where $g \in F - H$ and $m > 1$. Then there is $f \in F - H$ such that $f^m \in H$ and $|f| \leq [(2n)^k k^{2m} + 1](k+1) + 2k$.

Let f be the shortest element in $F - H$ such that $f^m \in H$. We write f as $f = ara^{-1}$ where r is cyclically reduced in X. Since $f^m = ar^m a^{-1} \in H$, every initial segment of the word $ar^m a^{-1}$ can be read as a label of a path in $\Gamma(H)$ originating at 1_H. For each $i = 0, 1, \dots, m$ let v_i be the terminal vertex of a path from 1_H in $\Gamma(H)$ labeled ar^i. Since $f = ara^{-1} \notin H$, we conclude that $v_0 \neq v_1$. Similarly, since $f^m = ar^m a^{-1} \in H$, we have $v_0 = v_m$. Thus for any word a' which is a label of a path in $\Gamma(H)$ from 1_H to v_0 we have $a' r (a')^{-1} \notin H$ and $a' r^m (a')^{-1} \in H$. Since f was assumed to be shortest in $F - H$ with $f^m \in H$, this implies that the path from 1_H to v_0 in $\Gamma(H)$ labeled a has no cycles and hence $|a| \leq k = \#V\Gamma$.

Let $r = b' c b''$ where c is the label of a simple loop in the path labeled r from v_0 to v_1 in $\Gamma(H)$ (so that $|c| \leq k$). Denote by u_i and w_i the terminal vertices of the paths labeled b' and $b'c$ originating from v_i. Thus c is the label of a path from u_i to w_i. Moreover, $u_0 = w_0$ by the choice of c. Consider the tuple $Q = (c, u_0, u_1, w_1, u_2, w_2, \dots, u_{m-1}, w_{m-1})$. Since $|c| \leq k$, there are at most $N := (2n)^k k^{2m}$ possibilities for Q.

Suppose $|r| \geq (N+1)(k+1)$. Since each path of length $k+1$ in $\Gamma(H)$ contains a cycle, we can write r as $r = b_0 c_0 b_1 c_1 \dots c_N b_{N+1}$ where each c_j has length at most k and is the label of a nontrivial simple loop in the path labeled r from v_0 to v_1 in $\Gamma(H)$. Let Q_j be the tuple produced as above for each $j = 0, 1, \dots, N$. By the choice of N there are some $j < s$ such that $Q_j = Q_s$. Let $z = c_j b_{j+1} \dots b_s$. Thus z is a subword of r which corresponds to a cycle in each path labeled r from v_i to v_{i+1} for $i = 0, 1 \dots, m-1$.

Let r' be the word obtained by deleting the subword z from r. That is $r' = b_0 c_0 \dots b_j c_s b_{s+1} \dots b_{N+1}$. By construction $|r'| < |r|$, $ar'a^{-1} \notin H$ and $a(r')^m a^{-1} \in H$. This contradicts the choice of f. Thus $|r| < (N+1)(k+1)$. Since $|a| \leq k$, this implies that $|f| = |ara^{-1}| \leq (N+1)(k+1) + 2k$ and the Claim holds.

Recall that by Proposition 7.10 if $g^m \in H$ then $g^{m_0} \in H$ for some $1 \leq m_0 \leq k = \#V\Gamma$. We can now formulate the algorithm for checking if H is isolated.

First, construct the graph $\Gamma = \Gamma(H)$ and compute $k = \#V\Gamma$. Put $M := [(2n)^k k^{2k} + 1](k+1) + 2k$. Now for each word f in $F(X)$ with $|f| \leq M$ such that $f \notin H$ check if there is $2 \leq m \leq k$ such that $f^m \in H$. If yes then H is not isolated. If no such f exists, then H is isolated. □

Remark 13.2. The above theorem is one of the few places in the present paper where our algorithm has high complexity. In fact, quite a few improvements are possible in the algorithm provided in Theorem 13.1. For example, if a subgroup is malnormal (which is easy to check) then it is also isolated. However, at present we do not know of a fast algorithm (comparable in speed with the malnormality test provided by Theorem 9.10) for determining if a finitely generated subgroup of a free group is isolated.

Lemma 13.3. *Let K be a finitely generated subgroup of $F(X)$ which is not malnormal in $F(X)$. Let $g \in F(X)$ be such that $g \notin K$ and $gKg^{-1} \cap K \neq 1$. Then the subgroup $H = \langle K, g \rangle$ is an algebraic extension of K. Moreover, the rank of H is less than or equal to the rank of K.*

Proof. Suppose, on the contrary, that $K \leq H$ is a free extension. Then $K \leq K'$ where $H = K' * C$ and $C \neq 1$. Note that $g^{-1} K' g \cap K' \neq 1$. On the other hand K' is a free factor of H and so malnormal in H. Hence $g \in K'$. Since $K \leq K', g \in K'$, we have $H = \langle K, g \rangle \leq K'$. This is impossible since $H = K' * C$ and $C \neq 1$. Thus $K \leq H$ is a algebraic extension of K, as required.

Since $H = \langle K, g \rangle$ and H, K are free, then either $rk(H) = rk(K) + 1$ and $H = K * \langle g \rangle$, or $rk(H) \leq rk(K)$. The former contradicts our assumption that $g^{-1} K g \cap K \neq 1$. Thus $rk(H) \leq rk(K)$, as required. □

Lemma 13.4. *Let K be a finitely generated subgroup of $F(X)$. Then there exists a unique subgroup H of $F(X)$ such that H is minimal among algebraic extensions of K which are malnormal in $F(X)$.*

Proof. If K is malnormal in $F(X)$, the statement is obvious. Assume now that K is not malnormal.

First we observe that there is at least one malnormal subgroup which is an algebraic extension of K.

Let $K = K_1$ and define a sequence $K_1 \leq K_2 \leq \dots$ as follows. Since K_1 is not malnormal, there is $g_2 \in F(X) - K_1$ such that $g_2 K_1 g_2^{-1} \cap K_1 \neq 1$. Put $K_2 = \langle K_1, g_2 \rangle$. Suppose now that $K_1 \leq K_2 \leq \dots \leq K_i$ are already constructed. If K_i is malnormal in $F(X)$, we terminate the sequence. Otherwise there is $g_{i+1} \in F(X) - K_i$ such that $g_{i+1} K_i g_{i+1}^{-1} \cap K_i \neq 1$. Put $K_{i+1} = \langle K_i, g_{i+1} \rangle$.

Claim. The sequence $K_1 \leq K_2 \leq \dots$ terminates in a finite number of steps. Moreover, each K_i is a algebraic extension of K and $rk(K_i) \leq rk(K)$.

Lemma 13.3 implies that $K_i \leq K_{i+1}$ is a algebraic extension and $rk(K_{i+1}) \leq rk(K_i)$. By transitivity (see Theorem 11.6) this means that $K \leq K_i$ is algebraic for each $i \geq 1$. Since $K_i \neq K_{i+1}$ and there are only finitely many algebraic extensions of K, the sequence terminates. Also, obviously $rk(K_i) \leq rk(K_1) = rk(K)$. Thus the Claim holds.

Suppose the sequence $K_1 \leq K_2 \leq \dots$ terminates in m steps with K_m. By construction this means that K_m is malnormal. The Claim also shows that $K \leq K_m$ is a algebraic extension.

Thus the set $\mathcal{M}(H)$ of algebraic extensions of K which are malnormal in $F(X)$ is nonempty. This sets is also obviously finite. We claim that it has a unique minimal element. Suppose this is not the case and H_1, H_2 are two such distinct minimal elements. Since H_1 and H_2 are malnormal in $F(X)$ and finitely generated, their intersection $H = H_1 \cap H_2$ is also malnormal in $F(X)$, finitely generated and contains K. Since $H_1 \neq H_2$ and H_1, H_2 are minimal in $\mathcal{M}(H)$, the group H is not an algebraic extension of K. That is the extension $K \leq H$ is free.

Let $C \neq 1, K'$ be such that $K \leq K'$, $H = K' * C$ and C is of the largest possible rank. By Lemma 11.9 K' is a algebraic extension of K. Since K' is a free factor in H, K' is malnormal in H and therefore in $F(X)$. Therefore $K' \in \mathcal{M}(H)$. However, $K' \neq H_1$, $K' \leq H_1$ which contradicts the minimality of H_1. □

Definition 13.5 (Malnormal closure)**.** Let $K \leq F(X)$ be a finitely generated subgroup. We denote

$$mal(K) := \cap\{H \mid K \leq H \leq F(X) \text{ and } H \text{ is malnormal in } F(X)\}$$

and refer to $mal(K)$ as the *malnormal closure* of K in $F(X)$.

Note that $F(X)$ is malnormal in $F(X)$, and hence $mal(K)$ is nonempty. Thus $mal(K)$ is the smallest malnormal subgroup of $F(X)$ containing K.

Theorem 13.6. *Let K be a finitely generated subgroup of $F(X)$.*

Let M be the unique minimal element among algebraic extensions of K which are malnormal in $F(X)$ (the existence of H follows from Lemma 13.4).

Then $M = mal(K)$. In particular $mal(K)$ is finitely generated and the extension $K \leq mal(K)$ is algebraic. Moreover, there is an algorithm which, given a finite set of generators of K, outputs a finite set of generators for $mal(K)$.

Proof. Let $K \leq H \leq F(X)$ where H is malnormal in $F(X)$. If H is an algebraic extension of K, then $M \leq H$ by the definition of M. Suppose now that the extension $K \leq H$ is free. Assume first that H is finitely generated. Let $C \neq 1$ and K' be such that $K \leq K', H = K' * C$ and C is of the largest possible rank. By Lemma 11.9 K' is an algebraic extension of K. Since K' is a free factor of H, K is malnormal in H and so in $F(X)$. Hence $M \leq K'$ by the definition of M and therefore $M \leq H$.

Assume now that H is not finitely generated, $K \leq H \leq F(X)$ and H is malnormal in $F(X)$. Thus H is a free group of infinite rank. Choose a free basis Y of H. Since K is finitely generated, the elements of a finite generating set of K involve only finitely many letters of Y. Thus K is contained in a finitely generated free factor H' of H. Since H' is a free factor of H, it is malnormal in H and so in $F(X)$. Hence $M \leq H'$ by the previous case. Therefore $M \leq H$.

Since H was chosen arbitrarily, this implies that

$$M \leq mal(K)$$

On the other hand $mal(K) \leq M$ by definition of $mal(K)$ since $K \leq M$ and M is malnormal in $F(X)$. Hence $M = mal(K)$, as required.

Thus $K \leq mal(K)$ is an algebraic extension and $mal(K)$ is finitely generated. Moreover, by Theorem 11.3 we can compute all algebraic extensions of K and for each of them check if it is malnormal in

$F(X)$ using Corollary 9.11. Since the number of algebraic extensions of K is finite, we can effectively find the unique minimal element among malnormal algebraic extensions of K, that is M. Thus, we can compute the malnormal closure $mal(K) = M$.

It remains to show that $rk(M) \leq rk(K)$. If $M = K$, there is nothing to prove. Suppose $K \lneq M$. Hence K is not malnormal in M by the choice of M. Therefore there is $g_1 \in M - K$ such that $g_1 K g_1^{-1} \cap K \neq 1$. Arguing exactly as in the proof of Lemma 13.4 we can find a sequence of elements $g_1, g_2, \dots inM$ such that $K = K_1 \lneq K_2 \lneq \dots$ and $K_{i+1} = \langle K_i, g_i \rangle$, where $g_i \in M - K_i$ is such that $g_i K_i g_i^{-1} \cap K_i \neq 1$. Again, by Lemma lem:step each $K_i \leq K_{i+1}$ is an algebraic extension and $rk(K_{i+1}) \leq rk(K_i)$. Hence each $K \leq K_i$ is an algebraic extension and $rk(K_i) \leq rk(K)$. Since the number of algebraic extensions of K in M is finite, any chain of subgroups of this sort terminates. Consider a maximal chain of this type $K = K_1 \lneq K_2 \lneq \cdots \lneq K_s$. Then K_s is malnormal in M. Indeed, if not then we can find $g_s \in M - K_s$ with $g_s K_s g_s^{-1} \cap K_s \neq 1$ and extend the chain, contradicting its maximality. Thus K_s is indeed malnormal in M and therefore in $F(X)$. Since K_s is an algebraic extension of K, we have $K_s = M$ by the definition of M. Therefore $rk(M) \leq rk(K)$, as required.

This completes the proof of Theorem 13.6. □

It turns out that one can use the same argument to analyze *isolators* of finitely generated subgroups of $F(X)$. Recall that a subgroup H of a group G is said to be *isolated* in G if whenever $g^n \in H$, we have $g \in H$. Note that if G is torsion-free and H is not isolated in G then H is not malnormal in G. Indeed, suppose $g^n \in H$ but $g \notin H$. Then $1 \neq \langle g^n \rangle = g \langle g^n \rangle g^{-1} \cap \langle g^n \rangle \leq gHg^{-1} \cap H$ and so H is not malnormal.

The proof of the following statement is a complete analogue of the proof of Lemma 13.4.

Lemma 13.7. *Let K be a finitely generated subgroup of $F(X)$. Then there exists a unique subgroup H of $F(X)$ such that H is minimal among algebraic extensions of K which are isolated in $F(X)$.*

Proof. If K is isolated in $F(X)$, the statement is obvious. Assume now that K is not isolated in $F(X)$ (and so not malnormal).

First we observe that there is at least one isolated subgroup which is an algebraic extension of K.

Let $K = K_1$ and define a sequence $K_1 \leq K_2 \leq \dots$ as follows. Since K_1 is not isolated, there is $g_2 \in F(X) - K_1$ such that $g_2^{n_2} \in K_1$. Put $K_2 = \langle K_1, g_2 \rangle$. Suppose now that $K_1 \leq K_2 \leq \cdots \leq K_i$ are already constructed. If K_i is isolated in $F(X)$, we terminate the sequence. Otherwise there is $g_{i+1} \in F(X) - K_i$ such that $g_{i+1}^{n_{i+1}} \in K_i$. Put $K_{i+1} = \langle K_i, g_{i+1} \rangle$.

Claim. The sequence $K_1 \leq K_2 \leq \dots$ terminates in a finite number of steps. Moreover, each K_i is a algebraic extension of K and $rk(K_i) \leq rk(K)$.

Note that $g_{i+1} K_i g_{i+1}^{-1} \cap K_i \neq 1$ and so Lemma 13.3 implies that $K_i \leq K_{i+1}$ is a algebraic extension and $rk(K_{i+1}) \leq rk(K_i)$. By transitivity (see Theorem 11.6) this means that $K \leq K_i$ is algebraic for each $i \geq 1$. Since $K_i \neq K_{i+1}$ and there are only finitely many algebraic extensions of K, the sequence terminates. Also, obviously $rk(K_i) \leq rk(K_1) = rk(K)$. Thus the Claim holds.

Suppose the sequence $K_1 \leq K_2 \leq \dots$ terminates in m steps with K_m. By construction this means that K_m is isolated in $F(X)$. The Claim also shows that $K \leq K_m$ is an algebraic extension.

Thus the set $\mathcal{I}(H)$ of algebraic extensions of K which are isolated in $F(X)$ is nonempty. This set is also obviously finite. We claim that it has a unique minimal element. Suppose this is not the case and H_1, H_2 are two such distinct minimal elements. Since H_1 and H_2 are isolated in $F(X)$ and finitely generated, their intersection $H = H_1 \cap H_2$ is also isolated in $F(X)$, finitely generated and contains K. Since $H_1 \neq H_2$ and H_1, H_2 are minimal in $\mathcal{I}(H)$, the group H is not an algebraic extension of K. That is the extension $K \leq H$ is free.

Let $C \neq 1, K'$ be such that $K \leq K'$, $H = K' * C$ and C is of the largest possible rank. By Lemma 11.9 K' is a algebraic extension of K. Since K' is a free factor in H, K' is isolated in H and hence in $F(X)$. Therefore $K' \in \mathcal{I}(H)$. However, $K' \neq H_1$, $K' \leq H_1$ which contradicts the minimality of H_1. □

Definition 13.8 (Isolator)**.** Let $K \leq F(X)$ be a finitely generated subgroup. We define the *isolator* $iso(K)$ of K in $F(X)$ as

$$iso(K) := \cap\{H \mid K \le H \le F(X) \text{ and } H \text{ is isolated in } F(X)\}.$$

Thus $iso(K)$ is the smallest isolated subgroup of $F(X)$ containing K.

Theorem 13.9. *Let K be a finitely generated subgroup of $F(X)$. Let M be the unique minimal element among algebraic extensions of K which are isolated in $F(X)$ (the existence of H follows from Lemma 13.7).*

Then $iso(K) = M$. In particular, $iso(K)$ is finitely generated, the extension $K \le iso(K)$ is algebraic and $rk(iso(K)) \le rk(K)$. Moreover, there exists an algorithm which, given a finite generating set of K, produces a free basis of $iso(K)$.

Proof. The proof is exactly the same as that of Theorem 13.6 (for the algorithmic part one needs to use Theorem 13.1 to decide if a particular algebraic extension of K is isolated in $F(X)$). We leave the details to the reader. □

14. Ascending chains of subgroups

It is well known that $F(X)$ contains free subgroups of infinite rank if $\#X \ge 2$. Therefore ascending chains of finitely generated subgroups of $F(X)$ do not necessarily terminate. However, by a classical result of M.Takahasi [46] and G.Higman [21] such chains do terminate if the ranks of the subgroups are bounded:

Theorem 14.1 (Takahasi-Higman)**.** [46, 21] *Let $F = F(X)$ be a free group of finite rank. Let $M \ge 1$ be an integer. Then every strictly ascending chain of subgroups of $F(X)$ of rank at most M terminates.*

Proof. Suppose the statement of Theorem 14.1 fails and there is an infinite strictly ascending chain of nontrivial subgroups of $F(X)$

$$K_1 \gneqq K_2 \gneqq \dots \tag{2}$$

where $rk(K_i) \le m$ for each $i \ge 1$.

Let $\Gamma_i = \Gamma(K_i)$ for $i = 1, 2, \dots$.

We will define the sequence of finite graphs $\Delta_1 \subseteq \Delta_2 \subseteq \dots$ as follows.

Step 1. Consider the graphs $\Gamma_{K_2}(K_1), \Gamma_{K_3}(K_1), \dots,$, which are all principal quotients of $\Gamma(K_1)$. Since $\Gamma(K_1)$ is finite, it has only finitely many principal quotients. Therefore for infinitely many values of $i \ge 2$ all $(\Gamma_{K_i}(K_1), 1_{K_i})$ are the same based X-digraph which we call (Δ_1, v_1).

Thus after passing to a subsequence in (2) we assume that $(\Gamma_{K_i}(K_1), 1_{K_i}) = (\Delta_1, v_1)$ for all $i \ge 2$.

Step n. Suppose the graphs $(\Delta_1, v_1) \gneqq (\Delta_2, v_2) \gneqq \cdots \gneqq (\Delta_{n-1}, v_{n-1})$ have already been constructed and the sequence (2) has been modified so that for each $j = 1, \dots, n-1$ we have

$$\Gamma_{K_i}(K_j), 1_{K_i}) = (\Delta_j, v_j) \text{ for all } i \ge j+1.$$

Note that there exists $m \ge n$ such that $\Gamma_{K_m}(K_{n-1}), 1_{K_m}) = (\Delta_{n-1}, v_{n-1})$ is a proper subgraph of $(\Gamma(K_m), 1_{K_m})$. If this is not the case then the sequence of graphs $(\Gamma(K_i), 1_{K_i})$ eventually stabilizes and so $K_i = K_{i+1}$ for some i, contrary to our assumptions.

After passing to a subsequence in (2) we may assume that $(\Gamma_{K_n}(K_{n-1}), 1_{K_n}) = (\Delta_{n-1}, v_{n-1})$ is a proper subgraph of $(\Gamma(K_n), 1_{K_n})$.

Since $(\Gamma(K_n), 1_{K_n})$ is a finite graph and has only finitely many principal quotients, there are infinitely many indices $i \ge n+1$ such that all $(\Gamma_{K_i}(K_n), 1_{K_i})$ are the same based X-digraph which we call (Δ_n, v_n). Once again, after passing to a subsequence in (2) we may assume that $(\Gamma_{K_i}(K_n), 1_{K_i}) = (\Delta_n, v_n)$ for all $i \ge n+1$.

Note that (Δ_{n-1}, v_{n-1}) is a subgraph of (Δ_n, v_n). Indeed, the graph (Δ_{n-1}, v_{n-1}) can be obtained by first mapping $(\Gamma(K_{n-1}), 1_{K_{n-1}})$ into $(\Gamma(K_n), 1_{K_n})$ (under the graph morphism corresponding to $K_{n-1} \le K_n$) and then mapping the result into $(\Gamma(K_{n+1}), 1_{K_{n+1}})$ (under the graph morphism corresponding to $K_n \le K_{n+1}$). The full image of the last map from $(\Gamma(K_{n-1}), 1_{K_{n-1}})$ into $(\Gamma(K_{n+1}), 1_{K_{n+1}})$ is precisely

(Δ_n, v_n). Thus (Δ_{n-1}, v_{n-1}) is indeed a subgraph of (Δ_n, v_n). Moreover it is a proper subgraph. Indeed, suppose $(\Delta_{n-1}, v_{n-1}) = (\Delta_n, v_n)$. Since $(\Gamma_{K_{n+1}}(K_n), 1_{K_{n+1}}) = (\Delta_n, v_n) = (\Delta_{n-1}, v_{n-1})$, we have $K_n \le L(\Delta_{n-1}, v_{n-1})$. On the other hand by our assumption (Δ_{n-1}, v_{n-1}) is a proper subgraph of $\Gamma(K_n), 1_{K_n})$. Since Δ_{n-1} is a core graph with respect to v_{n-1}, this implies that $L(\Delta_{n-1}, v_{n-1}) \lneqq K_n$, yielding a contradiction. Thus (Δ_{n-1}, v_{n-1}) is a proper subgraph of (Δ_n, v_n). We now go to the next step.

This procedure gives us an infinite sequence of core graphs $(\Delta_1, v_1) \subsetneqq (\Delta_2, v_2) \subsetneqq \dots$ such that for each $n \ge 1$ (Δ_n, v_n) is a subgraph of one of the graphs $(\Gamma(K_i), 1_{K_i})$ for some group K_i from the *original* sequence (2). However, since (Δ_{n-1}, v_{n-1}) is a proper subgraph of (Δ_n, v_n), the group $L(\Delta_{n-1}, v_{n-1})$ is a proper free factor of the group $L(\Delta_n, v_n)$ for every $n \ge 1$. Therefore $rk(L(\Delta_n, v_n)) \ge n$ for each $n \ge 1$. However, (Δ_{M+1}, v_{M+1}) is a subgraph of some $(\Gamma(K_i), 1_{K_i})$ and therefore $L(\Delta_{M+1}, v_{M+1})$ is a free factor of K_i. Since the rank of $L(\Delta_{M+1}, v_{M+1})$ is at least $M+1$, this implies that the rank of K_i is at least $M+1$ as well. This contradicts our assumption in Theorem 14.1 that the ranks of all subgroups in (2) are at most M. ☐

15. Descending chains of subgroups

In this section we will show that graph-theoretic methods developed earlier allow one to work with descending as well as ascending chains of subgroups. Namely, we will prove two classical results due to M.Takahasi [46].

Theorem 15.1. [46] *Suppose K_i are subgroups of $F(X)$ which form an infinite strictly descending chain*

$$K_1 \gneqq K_2 \gneqq \dots$$

Put $K_\infty = \cap_{n=1}^{\infty} K_n$. Let K be a finitely generated free factor of K_∞. Then K is a free factor in all but finitely many K_i.

Proof. Note that $\Gamma(K)$ is a finite graph and therefore it has only finitely many principal quotients.

Therefore there is an infinite strictly increasing sequence of indices $(i_n)_{n=1}^{\infty}$ such that for all $n \ge 1$ the based graphs $(\Gamma_{K_{i_n}}(K), 1_{K_{i_n}})$ are the same graph (Δ, v). Let $H = L(\Delta, v)$ so that $(\Gamma(H), 1_H) = (\Delta, v)$.

We claim that in fact $(\Gamma_{K_j}(K), 1_{K_j}) = (\Delta, v)$ for each $j \ge i_1$. Indeed, suppose $i_n < j < i_{n+1}$. Since $i_n < j$ and $K_j \le K_{i_n}$ the graph $(\Gamma_{K_j}(K), 1_{K_j})$ is an epimorphic image of the graph $(\Gamma_{K_{i_n}}(K), 1_{K_{i_n}}) = (\Delta, v)$. Therefore $L(\Gamma_{K_j}(K), 1_{K_j}) \le L(\Delta, v)$. On the other hand $j < i_{n+1}$, $K_{i_{n+1}} \le K_j$ and therefore the graph $(\Delta, v) = (\Gamma_{K_{i_{n+1}}}(K), 1_{K_{i_{n+1}}})$ is an epimorphic image of the graph $(\Gamma_{K_j}(K), 1_{K_j})$. Hence $L(\Delta, v) \le L(\Gamma_{K_j}(K), 1_{K_j})$. Thus $L(\Gamma(H), 1_H) = L(\Delta, v) = L(\Gamma_{K_j}(K))$. Since Δ is a finite core graph with respect to v, this implies that $(\Gamma(H), 1_H) = (\Delta, v) = (\Gamma_{K_j}(K), 1_{K_j})$, as claimed.

Since $(\Gamma(H), 1_H) = (\Delta, v)$ is a subgraph of all $(\Gamma(K_j), 1_{K_j})$ for $j \ge i_1$, the group H is a free factor in all K_j for $j \ge i_1$. Therefore $H \le \cap_{i=1}^{\infty} K_i = K_\infty$.

Also, since (Δ, v) is an epimorphic image of $(\Gamma(K), 1_K)$, we have $K = L(\Gamma(K), 1_K) \le L(\Delta, v) = H$.

Recall that K is a finitely generated free factor of K_∞ and $K \le H \le K_\infty$. Therefore by Lemma 10.7 K is a free factor of H. Since H is a free factor of all but finitely many K_i, the same holds for K.

This completes the proof of Theorem 15.1 is proved. ☐

Theorem 15.2. [46] *Let $F(X)$ be a free group of finite rank. Let*

$$K_1 \gneqq K_2 \gneqq \dots$$

be an infinite strictly decreasing chain of subgroups in $F(X)$. Suppose also that $rk(K_i) \le M$ for each $i \ge 1$. Put $K_\infty = \cap_{i=1}^{\infty} K_i$.

Then:

(a) K_∞ is a free factor in all but finitely many K_i.

(b) The group K_∞ is finitely generated and its rank is at most $M-1$.

Proof. Note that (a) implies (b). In fact suppose (a) holds. Since K_∞ is a free factor in some K_i, the free group K_∞ is finitely generated and has rank at most M. Suppose that in fact K_∞ is free of rank M. Since a proper free factor of a free group always has small rank and $rk(K_i) \le M$ this implies that

$K_\infty = K_i$. However $K_{i+1} \subsetneqq K_i$ and therefore K_{i+1} is not contained in K_∞. This contradicts the fact that $K_\infty = \cap_{n=1}^{\infty} K_n$. Thus $rk(K_\infty) \le M - 1$ as required.

We will now establish part (a) of the theorem.

Suppose first that K_∞ is finitely generated. Since K_∞ is a free factor of itself, Theorem 15.1 implies that K_∞ is a free factor in all but finitely many K_i, as required.

Suppose now that K_∞ is a free group of infinite rank. Then there are finitely generated free factors of K_∞ of arbitrarily large rank. This contradicts the conclusion of Theorem 15.1 since the ranks of K_i are bounded. □

References

1. G. Baumslag, A. Myasnikov, and V. Remeslennikov, *Algebraic geometry over groups. I. Algebraic sets and ideal theory*, J. Algebra **219** (1999), no. 1, 16–79. MR 1 707 663
2. ———, *Malnormality is decidable in free groups*, Internat. J. Algebra Comput. **9** (1999), no. 6, 687–692. MR 1 727 165
3. M. Bestvina and M. Feighn, *Bounding the complexity of simplicial group actions on trees*, Invent. Math. **103** (1991), no. 3, 449–469. MR **92c:**20044
4. ———, *A combination theorem for negatively curved groups*, J. Differential Geom. **35** (1992), no. 1, 85–101. MR **93d:**53053
5. ———, *Addendum and correction to: "A combination theorem for negatively curved groups" [J. Differential Geom.* **35** *(1992), no. 1, 85–101; MR 93d:53053]*, J. Differential Geom. **43** (1996), no. 4, 783–788. MR **97h:**53045
6. M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. (2) **135** (1992), no. 1, 1–51. MR **92m:**20017
7. ———, *Train-tracks for surface homeomorphisms*, Topology **34** (1995), no. 1, 109–140. MR **96d:**57014
8. S. A. Bleiler and A. C. Jones, *The free product of groups with amalgamated subgroup malnormal in a single factor*, J. Pure Appl. Algebra **127** (1998), no. 2, 119–136. MR **99f:**20041
9. J. A. Bondy and U. S. R. Murty, *Graph theory with applications*, American Elsevier Publishing Co., Inc., New York, 1976. MR 54 #117
10. W. Dicks and E. Formanek, *The rank three case of the Hanna Neumann conjecture*, J. Group Theory **4** (2001), no. 2, 113–151. MR 1 812 321
11. W. Dicks and E. Ventura, *The group fixed by a family of injective endomorphisms of a free group*, American Mathematical Society, Providence, RI, 1996. MR **97h:**20030
12. Warren Dicks, *Equivalence of the strengthened Hanna Neumann conjecture and the amalgamated graph conjecture*, Invent. Math. **117** (1994), no. 3, 373–389. MR **95c:**20034
13. M. J. Dunwoody, *Folding sequences*, The Epstein birthday schrift, Geom. Topol., Coventry, 1998, pp. 139–158 (electronic). MR **2000f:**20037
14. ———, *A small unstable action on a tree*, Math. Res. Lett. **6** (1999), no. 5-6, 697–710. MR **2001c:**20051
15. S. M. Gersten, *Intersections of finitely generated subgroups of free groups and resolutions of graphs*, Invent. Math. **71** (1983), no. 3, 567–591. MR **85m:**05037b
16. ———, *The double exponential theorem for isodiametric and isoperimetric functions*, Internat. J. Algebra Comput. **1** (1991), no. 3, 321–327. MR **93b:**20058
17. R. Gitik, *On the combination theorem for negatively curved groups*, Internat. J. Algebra Comput. **6** (1996), no. 6, 751–760. MR **98b:**20058a
18. R. Z. Goldstein and E. C. Turner, *Monomorphisms of finitely generated free groups have finitely generated equalizers*, Invent. Math. **82** (1985), no. 2, 283–289. MR **88b:**57007
19. L. Greenberg, *Commensurable groups of Moebius transformations*, (1974), 227–237. Ann. of Math. Studies, No. 79. MR 52 #594
20. M. Hall, Jr., *A topology for free groups and related groups*, Ann. of Math. (2) **52** (1950), 127–139. MR 12,158b
21. G. Higman, *A finitely related group with an isomorphic proper factor group*, J. London Math. Soc. **26** (1951), 59–61. MR 12,390b
22. A. G. Howson, *On the intersection of finitely generated free groups*, J. London Math. Soc. **29** (1954), 428–434. MR 16,444c
23. I. Kapovich, *Howson property and one-relator groups*, Comm. Algebra **27** (1999), no. 3, 1057–1072. MR 1 669 108
24. ———, *A non-quasiconvexity embedding theorem for hyperbolic groups*, Math. Proc. Cambridge Philos. Soc. **127** (1999), no. 3, 461–486. MR 1 713 122
25. ———, *Mapping tori of endomorphisms of free groups*, Comm. Algebra **28** (2000), no. 6, 2895–2917. MR **2001c:**20098
26. I. Kapovich and R. Weidmann, *On the structure of two-generated hyperbolic groups*, Math. Z. **231** (1999), no. 4, 783–801. MR 1 709 496
27. I. Kapovich and D. T. Wise, *The equivalence of some residual properties of word-hyperbolic groups*, J. Algebra **223** (2000), no. 2, 562–583. MR **2001f:**20086

28. A. Karrass and D. Solitar, *The free product of two groups with a malnormal amalgamated subgroup*, Canad. J. Math. **23** (1971), 933–959. MR 47 #3541
29. O. Kharlampovich and A. Myasnikov, *Hyperbolic groups and free constructions*, Trans. Amer. Math. Soc. **350** (1998), no. 2, 571–613. MR **98d:**20041
30. ______, *Irreducible affine varieties over a free group. I. Irreducibility of quadratic equations and Nullstellensatz*, J. Algebra **200** (1998), no. 2, 472–516. MR **2000b:**20032a
31. ______, *Irreducible affine varieties over a free group. II. Systems in triangular quasi-quadratic form and description of residually free groups*, J. Algebra **200** (1998), no. 2, 517–570. MR **2000b:**20032b
32. R. C. Lyndon and P. E. Schupp, *Combinatorial group theory*, Springer-Verlag, Berlin, 1977, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89. MR 58 #28182
33. W. Magnus, A. Karrass, and D. Solitar, *Combinatorial group theory*, revised ed., Dover Publications Inc., New York, 1976, Presentations of groups in terms of generators and relations. MR 54 #10423
34. H. Neumann, *On the intersection of finitely generated free groups*, Publ. Math. Debrecen **4** (1956), 186–189. MR 18,11f
35. W. D. Neumann, *On intersections of finitely generated subgroups of free groups*, Groups—Canberra 1989, Springer, Berlin, 1990, pp. 161–170. MR **92b:**20026
36. J. Nielsen, *A basis for subgroups of free groups*, Math. Scand. **3** (1955), 31–43. MR 17,455e
37. ______, *A basis for subgroups of free groups*, Math. Scand. **3** (1955), 31–43. MR 17,455e
38. ______, *Jakob Nielsen: collected mathematical papers. Vol. 1*, Birkhäuser Boston Inc., Boston, MA, 1986, Edited and with a preface by Vagn Lundsgaard Hansen. MR **88a:**01070a
39. ______, *Jakob Nielsen: collected mathematical papers. Vol. 1*, Birkhäuser Boston Inc., Boston, Mass., 1986, Edited and with a preface by Vagn Lundsgaard Hansen. MR **88a:**01070a
40. ______, *Jakob Nielsen: collected mathematical papers. Vol. 2*, Birkhäuser Boston Inc., Boston, Mass., 1986, Edited and with a preface by Vagn Lundsgaard Hansen. MR **88a:**01070b
41. Ch. Schaufele and N. Zumoff, *$*$-groups, graphs, and bases*, Topology and combinatorial group theory (Hanover, NH, 1986/1987; Enfield, NH, 1988), Springer, Berlin, 1990, pp. 186–191. MR **92a:**20025
42. C. C. Sims, *Computation with finitely presented groups*, Cambridge University Press, Cambridge, 1994. MR **95f:**20053
43. J. R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565. MR **85m:**05037a
44. J. R. Stallings and A. R. Wolf, *The Todd-Coxeter process, using graphs*, Combinatorial group theory and topology (Alta, Utah, 1984), Princeton Univ. Press, Princeton, N.J., 1987, pp. 157–161. MR **88g:**20084
45. John R. Stallings, *Foldings of G-trees*, Arboreal group theory (Berkeley, CA, 1988), Springer, New York, 1991, pp. 355–368. MR **92g:**20040
46. M. Takahasi, *Note on chain conditions in free groups*, Osaka Math. J. **3** (1951), 221–225. MR 13,721e
47. R. Weidmann, *A Grushko theorem for 1-acylindrical splittings*, Bochum University, preprint, 1999.
48. D. T. Wise, *The residual finiteness of negatively curved polygons of finite groups*, Cornell University, preprint, 1997.
49. ______, *The residual finiteness of positive one-relator groups*, Cornell University, preprint, 1998.
50. ______, *A residually finite version of Rips' construction*, Cornell University, preprint, 1999.

Department of Mathematics, University of Illinois at Urbana-Champaign, 1409 West Green Street, Urbana, IL 61801, USA
E-mail address: `kapovich@math.uiuc.edu`

Department of Mathematics, City College of CUNY, New York, NY 10031, USA
E-mail address: `alexei@rio.sci.ccny.cuny.edu`